

Oracle® Solaris 11.1 Administration: Devices and File Systems

Copyright © 2004, 2013, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

Preface	15
1 Managing Removable Media (Tasks)	17
Managing Removable Media (Overview)	17
Removable Media Features and Benefits	17
Comparison of Manual and Automatic Mounting	18
Overview of Accessing Removable Media	18
Managing Removable Media	19
Removable Media Considerations	20
Formatting Diskettes	21
▼ How to Load Removable Media	21
▼ How to Format a Diskette (rmformat)	22
▼ How to Create a File System on Removable Media	23
▼ How to Create a File System on a DVD-RAM	24
▼ How to Check a File System on Removable Media	24
▼ How to Repair Bad Blocks on Removable Media	25
Applying Read or Write Protection and Password Protection to Removable Media	25
▼ How to Enable or Disable Write Protection on Removable Media	26
▼ How to Enable or Disable Read or Write Protection and Set a Password on Removable Media	26
Accessing Removable Media	27
Using Removable Media Names	27
Guidelines for Accessing Removable Media Data	27
▼ How to Add a New Removable Media Drive	28
▼ How to Disable or Enable Removable Media Services	28
▼ How to Access Information on Removable Media	29
▼ How to Determine If Removable Media Is Still in Use	29
▼ How to Eject Removable Media	30

Accessing Removable Media on a Remote System	30
▼ How to Make Local Media Available to Other Systems	31
▼ How to Access Removable Media on Remote Systems	32
2 Writing CDs and DVDs (Tasks)	35
Working With Audio CDs and Data CDs and DVDs	35
CD/DVD Media Commonly Used Terms	36
Writing CD and DVD Data and Audio CDs	37
Restricting User Access to Removable Media With RBAC	38
▼ How to Restrict User Access to Removable Media With RBAC	39
▼ How to Identify a CD or DVD Writer	39
▼ How to Check the CD or DVD Media	40
Creating a Data CD or DVD	40
▼ How to Create an ISO 9660 File System for a Data CD or DVD	41
▼ How to Create a Multi-Session Data CD	41
Creating an Audio CD	43
▼ How to Create an Audio CD	44
▼ How to Extract an Audio Track on a CD	45
▼ How to Copy a CD	45
▼ How to Erase CD-RW Media	46
3 Managing Devices (Tasks)	47
What's New in Device Management?	47
Support for USB 3.0 Devices	47
Where to Find Additional Device Management Tasks	48
Managing Devices in Oracle Solaris	48
x86: Identifying Device Support	48
About Device Drivers	49
Automatic Configuration of Devices	50
Displaying Device Configuration Information	51
Resolving Faulty Devices	55
Adding a Peripheral Device to a System	58
▼ How to Add a Peripheral Device	58
▼ How to Add a Device Driver	59
Accessing Devices	60

How Device Information Is Created	60
How Devices Are Managed	60
Device Naming Conventions	61
Logical Disk Device Names	61
Logical Tape Device Names	64
Logical Removable Media Device Names	64
4 Dynamically Configuring Devices (Tasks)	65
Dynamic Reconfiguration and Hot-Plugging	65
Attachment Points	67
Detaching PCI or PCIe Adapter Cards	69
Attaching PCI or PCIe Adapter Cards	69
PCIe Hot-Plugging With the (hotplug) Command	69
SCSI Hot-Plugging With the <code>cfgadm</code> Command (Task Map)	71
SCSI Hot-Plugging With the <code>cfgadm</code> Command	72
▼ How to Display Information About SCSI Devices	72
▼ How to Unconfigure a SCSI Controller	73
▼ How to Configure a SCSI Controller	74
▼ How to Configure a SCSI Device	74
▼ How to Disconnect a SCSI Controller	75
▼ SPARC: How to Connect a SCSI Controller	76
▼ SPARC: How to Add a SCSI Device to a SCSI Bus	76
▼ SPARC: How to Replace an Identical Device on a SCSI Controller	77
▼ SPARC: How to Remove a SCSI Device	79
Troubleshooting SCSI Configuration Problems	80
▼ How to Resolve a Failed SCSI Unconfigure Operation	81
PCI or PCIe Hot-Plugging With the <code>cfgadm</code> Command (Task Map)	81
PCI or PCIe Hot-Plugging With the <code>cfgadm</code> Command	82
PCIe LED Indicator Behavior	82
▼ How to Display PCI Slot Configuration Information	83
▼ How to Remove a PCI Adapter Card	84
▼ How to Add a PCI Adapter Card	86
Troubleshooting PCI Configuration Problems	87
SATA Hot-Plugging With the <code>cfgadm</code> Command	88
▼ How to Unconfigure a SATA Device	88

▼ How to Configure a SATA Device	89
Reconfiguration Coordination Manager (RCM) Script Overview	89
What Is an RCM Script?	90
What Can an RCM Script Do?	90
How Does the RCM Script Process Work?	90
RCM Script Tasks	91
Application Developer RCM Script (Task Map)	91
System Administrator RCM Script (Task Map)	92
Naming an RCM Script	93
Installing or Removing an RCM Script	93
▼ How to Install an RCM Script	93
▼ How to Remove an RCM Script	94
▼ How to Test an RCM Script	94
Tape Backup RCM Script Example	95
5 Managing USB Devices (Tasks)	99
What's New in USB Devices?	99
USB 3.0 Support	99
About USB Support in Oracle Solaris	100
Supported USB Features	100
USB Device Features and Compatibility Issues	101
Bus-Powered Devices	102
USB Keyboards and Mouse Devices	102
USB Host Controller and Hubs	103
USB Hub Devices	104
SPARC: USB Power Management	104
Guidelines for Connecting USB Cables	105
Overview of USB Devices	105
Commonly Used USB Acronyms	105
Oracle Solaris USB Architecture	106
USB Bus Description	107
USB Devices and Driver Classes	108
Managing USB Mass Storage Devices	109
Using USB Diskette Devices	110
Hot-Plugging USB Mass Storage Devices	111

Preparing to Use a USB Mass Storage Device	114
▼ How to Display USB Device Information	115
▼ How to Create a File System on a USB Mass Storage Device	116
▼ How to Modify Partitions and Create a PCFS File System on a USB Mass Storage Device	118
▼ How to Create a Solaris Partition and Modify the Slices on a USB Mass Storage Device ..	121
▼ How to Mount or Unmount a USB Mass Storage Device	123
Troubleshooting Tips for USB Mass Storage Devices	124
Disabling Specific USB Drivers	124
▼ How to Disable Specific USB Drivers	125
▼ How to Remove Unused USB Device Links	125
Using USB Audio Devices	126
Hot-Plugging Multiple USB Audio Devices	127
▼ How to Add USB Audio Devices	127
▼ How to Identify Your System's Primary Audio Device	128
▼ How to Change the Primary USB Audio Device	128
Troubleshooting USB Audio Device Problems	129
Hot-Plugging USB Devices With the <code>cfgadm</code> Command	129
▼ How to Display USB Bus Information (<code>cfgadm</code>)	130
▼ How to Unconfigure a USB Device	131
▼ How to Configure a USB Device	132
▼ How to Logically Disconnect a USB Device	132
▼ How to Logically Connect a USB Device	133
▼ How to Logically Disconnect a USB Device Subtree	133
▼ How to Reset a USB Device	134
▼ How to Change the Default Configuration of a Multi-Configuration USB Device	134
6 Using InfiniBand Devices (Overview/Tasks)	137
Overview of InfiniBand Devices	137
InfiniBand Software Packages	139
Dynamically Reconfiguring IB Devices (Task Map)	139
Dynamically Reconfiguring IB Devices (<code>cfgadm</code>)	141
Unconfiguring IB Device Considerations	142
▼ How to Display IB Device Information	142
▼ How to Unconfigure an IB Port, HCA_SVC, or a VPPA Device	144

▼ How to Configure a IB Port, HCA_SVC, or a VPPA Device	144
▼ How to Unconfigure an IB Pseudo Device	145
▼ How to Configure an IB Pseudo Device	145
▼ How to Display Kernel IB Clients of an HCA	146
▼ How to Dynamically Reconfigure an HCA With Active EoIB Devices	146
▼ How to Reconfigure and Restore an EoIB Interface After Hot Removal	147
Configuring an IB HCA	148
▼ How to Update the IB P_key Tables	148
▼ How to Display IB Communication Services	149
▼ How to Add a VPPA Communication Service	149
▼ How to Remove an Existing IB Port, HCA_SVC, or a VPPA Communication Service	150
Using the uDAPL Application Interface With InfiniBand Devices	150
▼ How to Enable uDAPL	151
Updating the DAT Static Registry	151
Administering IPoIB Devices (dladm)	152
▼ How to Display Physical Data Link Information	153
▼ How to Create IB Partition Links	153
▼ How to Display IB Partition Link Information	154
▼ How to Remove an IB Partition Link	155
Monitoring and Troubleshooting IB Devices	155
7 Managing Disks (Overview)	159
What's New in Disk Management?	159
Support for Booting From EFI (GPT) Labeled Disks	159
Installation Support on Large Disks	160
Advanced Format Disk Support	161
Where to Find Disk Management Tasks	162
Overview of Disk Management	162
Disk Terminology	162
About Disk Labels	163
EFI (GPT) Disk Label	163
About Disk Slices	167
format Utility	168
Partitioning a Disk	171
Partition Table Terminology	171

Displaying Partition Table Information	172
Using the Free Hog Slice	174
8 Managing Disk Use (Tasks)	175
Managing Disk Use (Task Map)	175
Displaying Information About Files and Disk Space	176
▼ How to Display Information About Files and Disk Space	177
Checking the Size of Files	178
▼ How to Display the Size of Files	179
▼ How to Find Large Files	180
▼ How to Find Files That Exceed a Specified Size Limit	181
Checking the Size of Directories	182
▼ How to Display the Size of Directories, Subdirectories, and Files	182
▼ How to Display the User Ownership of Local UFS File Systems	183
Finding and Removing Old or Inactive Files	184
▼ How to List the Newest Files	184
▼ How to Find and Remove Old or Inactive Files	185
▼ How to Clear Out Temporary Directories	186
▼ How to Find and Delete core Files	187
▼ How to Delete Crash Dump Files	187
9 Administering Disks (Tasks)	189
Administering Disks (Task Map)	189
Identifying Disks on a System	190
▼ How to Identify the Disks on a System	190
Formatting a Disk	192
▼ How to Determine if a Disk Is Formatted	192
▼ How to Format a Disk	193
Displaying Disk Slices	194
▼ How to Display Disk Slice Information	195
Creating and Examining a Disk Label	197
▼ How to Label a Disk	197
▼ How to Examine a Disk Label	202
Recovering a Corrupted Disk Label	203
▼ How to Recover a Corrupted Disk Label	204

Adding a Third-Party Disk	206
10 Setting Up Disks (Tasks)	207
SPARC: Setting up Disks (Task Map)	207
SPARC: Setting Up Disks for ZFS File Systems	208
▼ SPARC: How to Set Up a Disk for a ZFS Root File System	209
SPARC: Creating a Disk Slice for a ZFS Root File System	209
▼ SPARC: How to Create a Disk Slice for a ZFS Root File System	210
▼ SPARC: How to Install Boot Blocks for a ZFS Root File System	214
▼ SPARC: How to Set Up a Disk for a ZFS Non-Root File System	214
x86: Setting Up Disks for ZFS File Systems (Task Map)	215
x86: Setting Up Disks for ZFS File Systems	216
x86: Creating and Changing Solaris <code>fdisk</code> Partitions	229
x86: Guidelines for Creating an <code>fdisk</code> Partition	229
▼ x86: How to Create a Solaris <code>fdisk</code> Partition	230
Changing the <code>fdisk</code> Partition Identifier	232
▼ How to Change the Solaris <code>fdisk</code> Identifier	233
11 Configuring Storage Devices With COMSTAR (Tasks)	235
Overview of COMSTAR Features	235
Identifying COMSTAR Software and Hardware Requirements	236
Configuring Storage Devices With COMSTAR (Task Map)	236
COMSTAR Terminology	237
Configuring Dynamic or Static Target Discovery	237
Configuring Storage Devices With COMSTAR (Tasks)	238
▼ How to Enable the STMF Service	239
▼ How to Back Up and Restore a COMSTAR Configuration	240
▼ How to Create a Logical Unit	240
▼ How to Create an iSCSI Target	241
▼ How to Enable iSNS Discovery for the Target Device	242
▼ How to Configure an IB HCA for iSER	242
Creating iSCSI Target Portal Groups	243
▼ How to Access iSCSI Disks	245
Making SCSI Logical Units Available	245
▼ How to Make a Logical Unit Available to All Systems	246

▼ How to Restrict Logical Unit Access to Selected Systems	247
Configuring Fibre Channel Devices With COMSTAR	248
Configuring Fibre Channel Ports With COMSTAR	248
Making Logical Units Available for FC and FCoE	251
Configuring FCoE Devices With COMSTAR	252
Configuring FCoE Ports	252
Enabling 802.3x PAUSE and Jumbo Frames on the Ethernet Interface	253
▼ How to Create FCoE Target Ports	253
▼ How to Verify That an FCoE Target Port Is Working	254
▼ How to Delete FCoE Target Ports	254
Configuring SRP Devices With COMSTAR	255
Using COMSTAR Views With SRP	255
▼ How to Enable the SRP Target Service	256
▼ How to Verify SRP Target Status	256
12 Configuring and Managing the Oracle Solaris Internet Storage Name Service (iSNS)	257
The iSNS Technology (Overview)	257
Configuring the iSNS Server	259
Setting Up the iSNS Administrative Settings	260
Using the Command Line Interface to Configure iSNS	262
Managing the iSNS Server and Clients	265
▼ How to Display the Status of a Discovery Domain Set	266
▼ How to Display the Status of a Discovery Domain	266
▼ How to Display the Status of Clients	266
▼ How to Remove a Client from a Discovery Domain	267
▼ How to Remove a Discovery Domain from a Discovery Domain Set	267
▼ How to Disable a Discovery Domain Set	268
▼ How to Remove a Discovery Domain Set	268
13 The format Utility (Reference)	269
Recommendations and Requirements for Using the format Utility	269
format Menu and Command Descriptions	270
partition Menu	272
x86: fdisk Menu	272
analyze Menu	273

defect Menu	275
Rules for Input to format Commands	276
Specifying Numbers to format Commands	276
Specifying format Command Names	276
Specifying Disk Names to format Commands	277
Getting Help on the format Utility	277
14 Managing File Systems (Overview)	279
What's New in Oracle Solaris File Systems?	279
Shared File Systems Across Boot Environments	279
Where to Find File System Management Tasks	280
Overview of File Systems	280
Types of Oracle Solaris File Systems	280
Default Oracle Solaris File Systems	285
Overview of Mounting and Unmounting File Systems	287
The Mounted File System Table	287
The Virtual File System Table	288
The NFS Environment	288
Automounting (autofs)	289
The Oracle Solaris SMB Service	290
Determining a File System's Type	291
How to Determine a File System's Type	291
15 Creating and Mounting File Systems (Tasks)	293
Creating Oracle Solaris File Systems	293
Creating ZFS File Systems	293
Creating a Temporary File System	294
Creating a LOFS File System	294
Mounting and Unmounting Oracle Solaris File Systems	294
Field Descriptions for the <code>/etc/vfstab</code> File	296
Prerequisites for Unmounting Oracle Solaris File Systems	297
Creating and Mounting Oracle Solaris File Systems	298
▼ How to Create an ZFS File System	298
▼ How to Create and Mount a Legacy UFS File System	299
▼ How to Create and Mount a TMPFS File System	300

▼ How to Create and Mount an LOFS File System	301
▼ How to Add an Entry to the /etc/vfstab File	302
▼ How to Mount a File System (/etc/vfstab File)	303
▼ How to Mount an NFS File System (mount Command)	304
▼ x86: How to Mount a PCFS (DOS) File System From a Hard Disk (mount Command) ...	305
▼ How to Stop All Processes That Are Accessing a File System	306
▼ How to Unmount a File System	307
16 Configuring Additional Swap Space (Tasks)	309
About Swap Space	309
Swap Space and Virtual Memory	309
Swap Space and the TMPFS File System	310
Swap Space and Dump Device Configuration	311
Swap Space and Dynamic Reconfiguration	311
Configuring Swap Space in a SAN Environment	312
How Do I Know If I Need More Swap Space?	312
Swap-Related Error Messages	312
TMPFS-Related Error Messages	313
How Swap Space Is Allocated	313
Swap Areas and the /etc/vfstab File	313
Planning for Swap Space	314
Allocating Swap Space for ZFS-Based Systems	315
Monitoring Swap Resources	315
Adding or Changing Swap Space in an Oracle Solaris ZFS Root Environment	316
▼ How to Add Swap Space in an Oracle Solaris ZFS Root Environment	316
17 Copying Files and File Systems (Tasks)	319
Commands for Copying File Systems	319
Copying Directories Between File Systems (cpio Command)	321
▼ How to Copy Directories Between File Systems (cpio)	322
Copying Files and File Systems to Tape	323
Copying Files to Tape (tar Command)	323
▼ How to Copy Files to a Tape (tar)	323
▼ How to List the Files on a Tape (tar)	324
▼ How to Retrieve Files From a Tape (tar)	325

Copying Files to a Tape With the pax Command	326
▼ How to Copy Files to a Tape (pax)	326
Copying Files to Tape With the cpio Command	327
▼ How to Copy All Files in a Directory to a Tape (cpio)	327
▼ How to List the Files on a Tape (cpio)	328
▼ How to Retrieve All Files From a Tape (cpio)	328
▼ How to Retrieve Specific Files From a Tape (cpio)	329
Copying Files to a Remote Tape Device	330
▼ How to Copy Files to a Remote Tape Device (tar and dd)	330
▼ How to Extract Files From a Remote Tape Device	331
18 Managing Tape Drives (Tasks)	333
Choosing Which Media to Use	333
Backup Device Names	334
Specifying the Rewind Option for a Tape Drive	335
Specifying Different Densities for a Tape Drive	335
Displaying Tape Drive Status	336
▼ How to Display Tape Drive Status	336
Handling Magnetic Tape Cartridges	337
Retensioning a Magnetic Tape Cartridge	337
Rewinding a Magnetic Tape Cartridge	337
Guidelines for Drive Maintenance and Media Handling	337
Index	339

Preface

Oracle Solaris 11.1 System Administration Guide: Devices and File Systems is part of a set that includes a significant part of the Oracle Solaris system administration information. This guide contains information for both SPARC based and x86 based systems.

This book assumes you have completed the following tasks:

- Installed the Oracle Solaris software
- Set up all the networking software that you plan to use

New Oracle Solaris features of interest to system administrators are covered in sections called *What's New in ... ?* in the appropriate chapters.

Note – This Oracle Solaris release supports systems that use the SPARC and x86 families of processor architectures. The supported systems appear in the *Oracle Solaris Hardware Compatibility List* at <http://www.oracle.com/webfolder/technetwork/hcl/index.html>. This document cites any implementation differences between the platform types.

Who Should Use This Book

This book is intended for anyone responsible for administering one or more systems running the Oracle Solaris 11 release. To use this book, you should have 1–2 years of UNIX system administration experience. Attending UNIX system administration training courses might be helpful.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

What Typographic Conventions Mean

The following table describes the typographic conventions used in this book.

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
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. <code>machine_name% you have mail.</code>
AaBbCc123	What you type, contrasted with on screen computer output	<code>machine_name% su</code> Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <code>rm filename</code> .
<i>AaBbCc123</i>	Book titles, new words or terms, or words to be emphasized	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. Do <i>not</i> save changes yet.

Shell Prompts in Command Examples

The following table shows the default system prompt and root (or administrative) prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell prompt	<code>machine_name%</code>
C shell root prompt	<code>machine_name#</code>
Bourne shell and Korn shell prompt	<code>\$</code>
Bourne shell and Korn shell root prompt	<code>#</code>

Managing Removable Media (Tasks)

This chapter describes how to manage and access removable media in the Oracle Solaris OS.

This is a list of the information in this chapter:

- “Managing Removable Media (Overview)” on page 17
- “Managing Removable Media” on page 19
- “Accessing Removable Media” on page 27
- “Accessing Removable Media on a Remote System” on page 30

Managing Removable Media (Overview)

The following information is described in this section:

- “Removable Media Features and Benefits” on page 17
- “Comparison of Manual and Automatic Mounting” on page 18
- “Overview of Accessing Removable Media” on page 18

Removable Media Features and Benefits

The Oracle Solaris release gives users and software developers a standard interface for dealing with removable media. Removable media services provide the following benefits:

- Automatically mounts removable media. For a comparison of manual and automatic mounting, see the following section.
- Enables you to access removable media without having to become an administrator.
- Allows you to give other systems on the network automatic access to any removable media on your local system. For more information, see “[Accessing Removable Media on a Remote System](#)” on page 30.

Comparison of Manual and Automatic Mounting

The following table compares the steps involved in manual mounting (without removable media services) and automatic mounting (with removable media management) of removable media.

TABLE 1-1 Comparison of Manual and Automatic Mounting of Removable Media

Steps	Manual Mounting	Automatic Mounting
1	Insert media.	Insert media.
2	Become an administrator.	For USB diskettes, use the <code>volcheck</code> command.
3	Determine the location of the media device.	Removable media services automatically perform many of the tasks that are required to manually mount and work with removable media.
4	Create a mount point.	
5	Make sure you are not in the mount point directory.	
6	Mount the device and use the proper mount options.	
7	Exit the administrator account.	
8	Work with files on media.	Work with files on media.
9	Become an administrator.	
10	Unmount the media device.	
11	Eject media.	Eject media.
12	Exit the administrator account.	

Overview of Accessing Removable Media

Essentially, removable media services enable you to access removable media just as manual mounting does, but more easily and without the need for administrative access.

If the media contains a file system and a label, the media label name is used to name the `/media/pathname` mount point. If a label is not present, the disk model name is used to name the media, such as `/media/cdrom`. A generic *nickname* is used only for legacy symbolic links. For example, `/rmdisk/rmdisk0`.

If your system has more than one type of removable device, see the following table for their access points.

TABLE 1-2 How to Access Data on Removable Media

Access	Insert	Unlabeled Media Pathnames	Labeled Media Pathname Examples
Files on a removable hard disk	The removable hard disk and type <code>volcheck</code> on the command line	<code>/media/usb-disk</code> or the legacy path <code>/rmdisk/rmdisk0</code>	<code>/media/00JB-00CRA0</code>
Files on a DVD	The DVD and wait for a few seconds	<code>/media/cdrom</code>	<code>/media/Oracle_Solaris-11_1-Live-X86</code>

You can use the `rmmount -l` command to identify mounted media on your system. For example:

```
# rmmount -l
/dev/dsk/c5t0d0p0      rmdisk6,/media/FD-05PUB
/dev/dsk/c4t0d3p0      rmdisk5,/media/223UHS-SD-MMC
/dev/dsk/c2t0d0s2      cdrom1,cd1,sr1,Oracle_Solaris-11_1-AI-SPARC,/media/Oracle_Solaris-11_1-AI-SPARC
/dev/dsk/c3t0d0p0      rmdisk2,/media/00JB-00CRA0
```

In the above output, the mounted devices are as follows:

```
/dev/dsk/c5t0d0p0      USB diskette
/dev/dsk/c4t0d3p0      CF card in a USB card reader
/dev/dsk/c2t0d0s2      DVD-ROM
/dev/dsk/c3t0d0p0      Removable USB disk
```

Managing Removable Media

The following guidelines and procedures are described in this section:

- “Removable Media Considerations” on page 20
- “How to Load Removable Media” on page 21
- “How to Format a Diskette (`rmformat`)” on page 22
- “How to Create a File System on Removable Media” on page 23
- “How to Create a File System on a DVD-RAM” on page 24
- “How to Check a File System on Removable Media” on page 24
- “How to Repair Bad Blocks on Removable Media” on page 25
- “Applying Read or Write Protection and Password Protection to Removable Media” on page 25
- “How to Enable or Disable Write Protection on Removable Media” on page 26

- [“How to Enable or Disable Read or Write Protection and Set a Password on Removable Media” on page 26](#)

Removable Media Considerations

Keep the following considerations in mind when working with diskettes:

- File system formats in Oracle Solaris consist of the basic “bit” formatting, in addition to the structure to support an Oracle Solaris file system. A complete format for a DOS file system consists of the basic “bit” formatting in addition to the structure to support either an MS-DOS or an NEC-DOS file system. The procedures required to prepare a media for each type of file system are different. Therefore, before you format a diskette, consider which procedure to follow. For more information, see [“Managing Removable Media” on page 19](#).
 - For information on removable media names, see [“Using Removable Media Names” on page 27](#).
 - Diskettes that are not named (that is, they have no “label”) are assigned the default name of `unnamed_floppy`.
 - Diskettes that are not named (that is, they have no “label”) are assigned the default name of `floppy`.

An Oracle Solaris system can format the following file system types:

- ZFS or UFS
- MS-DOS or NEC-DOS (PCFS)
- UDFS

On an Oracle Solaris system (either SPARC or x86), you can format diskettes with the following densities.

Diskette Size	Diskette Density	Capacity
3.5”	High density (HD)	1.44 MB
3.5”	Double density (DD)	720 KB

By default, the diskette drive formats a diskette to a like density. This default means that a 1.44 MB drive attempts to format a diskette for 1.44 MB, regardless of whether the diskette is, in fact, a 1.44 MB diskette, unless you instruct it otherwise. In other words, a diskette can be formatted to its capacity or lower, and a drive can format to its capacity or lower.

Formatting Diskettes

You can use the `rmformat` command to format and perform other management tasks on removable media. File systems are mounted automatically. So, you might have to unmount media before you can format it, if the media contains an existing file system.

The `rmformat` command has three formatting options:

- `quick` – This option formats diskettes without certification or with limited certification of certain tracks on the media.
- `long` – This option completely formats diskettes. For some devices, the use of this option might include the certification of the whole media by the drive.
- `force` – This option formats completely without user confirmation. For diskettes with a password-protection mechanism, this option clears the password before formatting. This feature is useful when a password is forgotten. On diskettes without password protection, this option forces a long format.

▼ How to Load Removable Media

For information about removable media hardware considerations, see [“Removable Media Considerations” on page 20](#).

1 Insert the media.

2 Ensure that the media is formatted.

If you aren't sure, insert the media and check the status messages in the system console window, as described in [Step 3](#). If you need to format the media, go to [“How to Format a Diskette \(`rmformat`\)” on page 22](#).

3 (Optional) Notify volume management if you are using a legacy, non-USB diskette device.

```
$ volcheck -v
```

Two status messages are possible:

```
media was found
```

Volume management detected the media and will attempt to mount it in the directory described in [“Using Removable Media Names” on page 27](#).

If the media is formatted properly, no error messages appear in the console.

If the media is not formatted, the “media was found” message is still displayed. However, error messages similar to the following appear in the system console window:

```
fd0: unformatted diskette or no diskette in the drive
```

```
fd0: read failed (40 1 0)
```

```
fd0: bad format
```

You must format the media before volume management can mount it. For more information, see [“How to Format a Diskette \(rmformat\)” on page 22](#).

```
no media was found
```

Volume management did not detect the media. Ensure that the media is inserted properly, and run `volcheck` again. If unsuccessful, check the media, which could be damaged. You can also try to mount the media manually.

4 Verify that the media was mounted by listing its contents.

For example, do the following for a diskette:

```
$ ls /media/floppy
lost+found myfiles
```

▼ How to Format a Diskette (rmformat)

You can use the `rmformat` command to format a diskette. By default, this command creates two partitions on the media: partition 0 and partition 2 (the whole media).

1 Verify that removable media service is running. If so, you can use the shorter nickname for the device name.

```
# svcs hal dbus rmvolmgr
STATE      STIME    FMRI
online     Apr_09   svc:/system/dbus:default
online     Apr_09   svc:/system/hal:default
online     Apr_09   svc:/system/filesystem/rmvolmgr:default
```

For information on restarting removable media services, see [“How to Disable or Enable Removable Media Services” on page 28](#). For information on identifying media device names, see [“Using Removable Media Names” on page 27](#).

2 Format the diskette.

```
$ rmformat -F [ quick | long | force ] device-name
```

See [“Formatting Diskettes” on page 21](#) for more information on `rmformat` formatting options.

If the `rmformat` output indicates bad blocks, see [“How to Repair Bad Blocks on Removable Media” on page 25](#).

3 (Optional) Label the diskette with an 8-character label.

```
$ rmformat -b label device-name
```

For information on creating a DOS label, see [mkfs_pcfs\(1M\)](#).

Example 1-1 Formatting a Diskette

This example shows how to format a diskette.

```
$ rmformat -F quick /dev/rdiskette
Formatting will erase all the data on disk.
Do you want to continue? (y/n) y
.....
```

▼ How to Create a File System on Removable Media**1 (Optional) Format a diskette, if necessary.**

To format a USB diskette, use syntax similar to the following:

```
$ rmformat -F long /dev/rdisk/c11t0d0p0
```

2 (Optional) Create an alternate Solaris partition table.

```
$ rmformat -s slice-file device-name
```

A sample slice file appears similar to the following:

```
slices: 0 = 0, 30MB, "wm", "home" :
         1 = 30MB, 51MB :
         2 = 0, 94MB, "wm", "backup" :
         6 = 81MB, 13MB
```

3 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights” in Oracle Solaris 11.1 Administration: Security Services](#).

4 Determine the appropriate file system type and select one of the following:

- Create a PCFS file system. For example:

```
# mkfs -F pcfs -o nofdisk,size=9800 /dev/rdisk/c11t0d0p0
```

- Create a UDFS file system. For example:

```
# mkfs -F udfs /dev/rdisk/c0t1d0p0
```

▼ How to Create a File System on a DVD-RAM

Use this procedure to create a file system on a DVD-RAM.

1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Create a file system on the DVD-RAM device.

For example, create a UDFS file system, as follows:

```
# mkfs -F udfs /dev/rdisk/c0t0d0s2
```

3 Mount the file system.

For example, mount a UDFS file system, as follows:

```
# mount -F udfs /dev/dsk/c0t0d0s2 /mnt
```

4 Verify that you can read or write to the file system.

5 When finished, eject the DVD-RAM.

▼ How to Check a File System on Removable Media

1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Identify the file system type and select one of the following:

■ Check a UDFS file system, as follows:

```
# fsck -F udfs device-name
```

■ Check a PCFS file system, as follows:

```
# fsck -F pcfs device-name
```

Example 1–2 Checking a PCFS File System on Removable Media

The following example shows how check the consistency of a PCFS file system on media.

```
# fsck -F pcfs /dev/rdisk/c0t4d0s2
** /dev/rdisk/c0t4d0s2
** Scanning file system meta-data
```

```
** Correcting any meta-data discrepancies
1457664 bytes.
0 bytes in bad sectors.
0 bytes in 0 directories.
0 bytes in 0 files.
1457664 bytes free.
512 bytes per allocation unit.
2847 total allocation units.
2847 available allocation units.
```

▼ How to Repair Bad Blocks on Removable Media

You can only use the `rmformat` command to verify, analyze, and repair bad sectors that are found during verification if the drive supports bad block management. Most USB memory sticks do not support bad block management.

If the drive supports bad block management, a best effort is made to rectify the bad block. If the bad block cannot be rectified despite the best effort mechanism, a message indicates the failure to repair the media.

1 Repair bad blocks on removable media.

```
$ rmformat -c block-numbers device-name
```

Supply the block number in decimal, octal, or hexadecimal format from a previous `rmformat` session.

2 Verify the media.

```
$ rmformat -V read device-name
```

Applying Read or Write Protection and Password Protection to Removable Media

You can apply read protection or write protection, and set a password, on removable media that support this feature.

▼ How to Enable or Disable Write Protection on Removable Media

- 1 Determine whether you want to enable or disable write protection and select one of the following:

- Enable write protection.

```
$ rmformat -w enable device-name
```

- Disable write protection.

```
$ rmformat -w disable device-name
```

- 2 Verify whether the media's write protection is enabled or disabled.

```
$ rmformat -p device-name
```

▼ How to Enable or Disable Read or Write Protection and Set a Password on Removable Media

You can apply a password with a maximum of 32 characters for removable media that support this feature.

You will receive a warning message if you attempt to apply a password on media that does not support this feature.

- 1 Determine whether you want to enable or disable read protection or write protection and set a password. Select one of the following:

- Enable read protection or write protection.

```
$ rmformat -W enable device-name
Please enter password (32 chars maximum): xxx
Please reenter password:
```

```
$ rmformat -R enable device-name
Please enter password (32 chars maximum): xxx
Please reenter password:
```

- Disable read protection or write protection and remove the password.

```
$ rmformat -W disable device-name
Please enter password (32 chars maximum): xxx
```

```
$ rmformat -R disable device-name
Please enter password (32 chars maximum): xxx
```

2 Verify whether the media's read protection or write protection is enabled or disabled.

```
$ rmformat -p device-name
```

Accessing Removable Media

You can access information on removable media with or without using volume management. For information on accessing information on removable media with GNOME's File Manager, see the GNOME desktop documentation.

This following procedures are described in this section:

- [“Using Removable Media Names” on page 27](#)
- [“Guidelines for Accessing Removable Media Data” on page 27](#)
- [“How to Add a New Removable Media Drive” on page 28](#)
- [“How to Disable or Enable Removable Media Services” on page 28](#)
- [“How to Access Information on Removable Media” on page 29](#)
- [“How to Determine If Removable Media Is Still in Use” on page 29](#)
- [“How to Eject Removable Media” on page 30](#)

Using Removable Media Names

Removable media is now mounted automatically in the `/media` directory. However, symbolic links to `/media` are provided from previous media mount points, `/cdrom` and `/rmdisk`, for compatibility purposes.

For example, a compact flash memory card (`/dev/dsk/c4d0p0:1`) is mounted as follows:

```
$ ls /media/memory-card-name
```

For example, a USB memory stick (`/dev/dsk/c3t0d0s0`) is mounted, as follows:

```
$ ls /media/U3
```

Guidelines for Accessing Removable Media Data

Most DVDs are formatted to the ISO 9660 standard, which is portable. So, DVDs can be mounted by volume management.

To accommodate possible different formats, a DVD is split into slices. Slices are similar in effect to partitions on hard disks. The 9660 portion is portable. If you are having trouble mounting a DVD, particularly if it is an installation DVD, make sure that its file system is appropriate for your system's architecture. For example, you can check the label on DVD.

▼ How to Add a New Removable Media Drive

Generally, most modern bus types support hot-plugging. This means you can insert a disk in an empty slot and the system recognizes it.

For more information about hot-plugging devices, see [Chapter 4, “Dynamically Configuring Devices \(Tasks\).”](#)

- 1 Become an administrator.**
- 2 Connect the new media drive.**
See your hardware handbook for specific instructions.
- 3 Confirm that the system sees the new media drive.**

```
# rmformat  
Looking for devices...
```

▼ How to Disable or Enable Removable Media Services

Occasionally, you might want to manage media without using removable media services. This section describes how to disable and enable removable media services.

Disabling these services means that you would have to mount all media manually by using the mount command.

- 1 Ensure that the media is not being used.**
If you are not sure whether you have found all users of the media, use the fuser command, see [“How to Determine If Removable Media Is Still in Use”](#) on page 29.
- 2 Become an administrator.**
- 3 Select one of the following:**

- You can disable some or all removable media features in this release:
 - To prevent volumes from mounting outside of user sessions, disable the `rmvolmgr` service. For example:

```
# svcadm disable rmvolmgr
```
 - To prevent any volume management, disable the `dbus`, `hal`, and `rmvolmgr` services.

```
# svcadm disable rmvolmgr  
# svcadm disable dbus  
# svcadm disable hal
```

Disabling these services means that you would have to mount all media manually by using the mount command.

- **Enable removable media services.**

```
# svcadm enable rnmvolmgr
# svcadm enable dbus
# svcadm enable hal
```

▼ How to Access Information on Removable Media

- 1 **Insert the media.**

The media is mounted after a few seconds.

- 2 **List the contents of the media. For example:**

```
# ls /media/Oracle_Solaris-11_1-AI-SPARC
auto_install      export            proc              solarismisc.zlib
bin               home             reconfigure      system
boot             jack             root              tmp
dev              mnt              sbin
devices          platform         solaris.zlib
```

- 3 **(Optional) Copy the file that you identified in the previous step.**

Example 1-3 Accessing Information on Removable Media

This example shows how to access information on a USB memory stick.

```
$ ls /media/usb-name
```

This example shows how to access information on a DVD.

```
$ ls /media
Oracle_Solaris-11_1-AI-SPARC cdrom
```

▼ How to Determine If Removable Media Is Still in Use

- 1 **Become an administrator.**
- 2 **Identify the processes that are accessing the media.**

```
# fuser -u /media
```

The -u displays the user of the media.

For more information, see [fuser\(1M\)](#).

3 (Optional) Kill the process accessing the media.

```
# fuser -u -k /media
```

The -k option kills the processes accessing the media.



Caution – Killing the processes that are accessing the media should only be used in emergency situations.

4 Verify that the process is gone.

```
# pgrep process-ID
```

▼ How to Eject Removable Media

1 Ensure that the media is not being used.

Remember, media is “being used” if a shell or an application is accessing any of its files or directories. If you are not sure whether you have found all users of a DVD (for example, a shell hidden behind a desktop tool might be accessing it), use the `fuser` command. See [“How to Determine If Removable Media Is Still in Use” on page 29](#).

2 Eject the media.

```
# eject media
```

For example, for a DVD, you would do the following:

```
# eject cdrom
```

For example, for a USB memory stick, you would do the following:

```
# eject rmdisk0
```

Tip – You can view the removable device name with the `eject -l` command.

Accessing Removable Media on a Remote System

Removable media can be shared with remote systems by using NFS. For additional information about using NFS, see [Managing Network File Systems in Oracle Solaris 11.1](#).

The following procedures are described in this section:

- [“How to Make Local Media Available to Other Systems” on page 31](#)
- [“How to Access Removable Media on Remote Systems” on page 32](#)

▼ How to Make Local Media Available to Other Systems

You can configure your system to share its media drives to make any media in those drives available to other systems. One exception is musical CDs. Once your media drives are shared, other systems can access the media they contain simply by mounting them. For instructions, see “How to Access Removable Media on Remote Systems” on page 32.

1 Become an administrator.

2 Confirm that the media is loaded.

3 Share the media.

For example, you might type a command similar to the following:

```
# share -F nfs -o ro /media/Oracle_Solaris-11_1-AI-SPARC
```

4 Determine whether the NFS server service is running.

```
# svcs *nfs*
```

The following output is returned from the `svcs` command if NFS server service is running:

```
online          14:28:43 svc:/network/nfs/server:default
```

5 Identify the NFS server status, and select one of the following:

- If the NFS server service is running, go to [Step 8](#).
- If the NFS server service is *not* running, go to the next step.

6 Start the NFS server service.

```
# svcadm enable network/nfs/server
```

7 Verify that the NFS daemons are running.

For example:

```
# svcs -p svc:/network/nfs/server:default
STATE      STIME    FMRI
online     Aug_30   svc:/network/nfs/server:default
           Aug_30   319 mountd
           Aug_30   323 nfsd
```

8 Verify that the media is available to other systems.

If the media is available, its share configuration is displayed.

```
# share
- /media/Oracle_Solaris-11_1-AI-SPARC sec=sys,ro ""
```

Example 1-4 Making Local DVDs Available to Other Systems

The following example shows how to make any local DVD available to other systems on the network.

```
# share -F nfs -o ro /media
# svcs *nfs*
# svcadm enable network/nfs/server
# svcs -p svc:/network/nfs/server:default
# share
- /media/Oracle_Solaris-11_1-AI-SPARC ro ""
```

▼ How to Access Removable Media on Remote Systems

You can access media on a remote system by manually mounting the media into your file system. Also, the remote system must have shared its media according to the instructions in [“How to Make Local Media Available to Other Systems”](#) on page 31.

1 Select an existing directory to serve as the mount point. Or create a mount point.

```
$ mkdir /directory
```

where */directory* is the name of the directory that you create to serve as a mount point for the remote system's DVD.

2 Find the name of the media that you want to mount.

```
$ showmount -e system-name
```

3 As an administrator, mount the media.

```
# mount -F nfs -o ro system-name:/media/media-name local-mount-point
```

system-name: Is the name of the system whose media you will mount.

media-name Is the name of the media you want to mount.

local-mount-point Is the local directory onto which you will mount the remote media.

4 Log out as an administrator.**5 Verify that the media has been mounted.**

```
$ ls /mnt
```

Example 1-5 Accessing DVDs or CDs on Remote Systems

The following example shows how to automatically access the remote DVD named `Oracle_Solaris-11_1-AI-SPARC` from the remote system `starbug` using `autofs`.

```
$ showmount -e starbug
export list for starbug:
/media/Oracle_Solaris-11_1-AI-SPARC (everyone)
$ ls /net/starbug/media/
Oracle_Solaris-11_1-AI-SPARC
```


Writing CDs and DVDs (Tasks)

This chapter provides step-by-step instructions for writing and copying data CDs and DVDs and audio CDs with the `cdwr` command.

This is a list of the information in this chapter:

- “Working With Audio CDs and Data CDs and DVDs” on page 35
- “Writing CD and DVD Data and Audio CDs” on page 37

Working With Audio CDs and Data CDs and DVDs

You can use the `cdwr` command to write file systems for CDs and DVDs in ISO 9660 format with Rock Ridge or Joliet extensions on CD-R, CD-RW, DVD-RW, or DVD+RW media devices.

You can use the `cdwr` command to perform the following tasks:

- Create data CDs and DVDs.
- Create audio CDs.
- Extract audio data from an audio CD.
- Copy CDs and DVDs.
- Erase CD-RW media.

The `cdwr` command is available in the following releases:

- Oracle Solaris 10 releases
- Oracle Solaris 11 release, `media/cdwr` package

For information on recommended CD-R or CD-RW devices, go to <http://www.oracle.com/us/sun/index.html>.

CD/DVD Media Commonly Used Terms

This section defines commonly used terms related to CD/DVD media.

Term	Description
CD-R	CD read media that can be written once and after that, can only be read from.
CD-RW	CD rewritable media that can be written to and erased. CD-RW media can only be read by CD-RW devices.
DVD-R	Digital video disk (recordable) that can be written once and after that, can only be read from. These devices have much larger capacity than CD-R media.
DVD+R	Digital video disk (recordable) that can be written once and after that, can only be read from. DVD+R devices have more complete error management system than DVD-R, which allows for more accurate burning to media, independent of the quality of the media.
DVD-RW	Digital video disk (rewritable) with storage capacity equal to a DVD-R. This media can be re-recorded by first erasing the entire disk.
DVD+RW	Digital video disk (random-access rewritable) with storage capacity equal to a DVD+R. This medium allows overwriting of individual blocks without erasing the entire disk.
DVD-RAM	Digital video disk (random access memory, rewritable) with circular rather than spiral tracks and hard sectoring.
ISO 9660	ISO, an acronym for Industry Standards Organization, is an organization that sets standards for computer storage formats. An ISO 9660 file system is a standard CD or DVD file system that enables you to read the same CD or DVD on any major computer platform. The standard, issued in 1988, was written by an industry group named High Sierra, named after the High Sierra Hotel in Nevada. Almost all computers with CD or DVD drives can read files from an ISO 9660 file system.
Joliet extensions	Adds Windows file system information.
Rock Ridge extensions	Adds UNIX file system information. (Rock Ridge is named after the town in the movie Blazing Saddles.) Note – These extensions are not mutually exclusive. You can specify both <code>mkisofs -R</code> and <code>-j</code> options for compatibility with both systems. (See <code>mkisofs(1M)</code> for details.)

Term	Description
MMC-compliant recorder	Acronym for Multi Media Command, which means these recorders comply with a common command set. Programs that can write to one MMC-compliant recorder should be able to write to all other recorders.
Red Book CDDA	Acronym for Compact Disc Digital Audio, which is an industry standard method for storing digital audio on compact discs. Also known by the term “Red Book” format. The official industry specification calls for one or more audio files sampled in 16-bit stereo sound at a sampling rate of 44.1 kilohertz (kHz).

Commonly used terms when writing to CD media are listed in the following table.

Term	Description
blanking	The process of erasing data from the CD-RW media.
session	A complete track with lead-in and lead-out information.
track	A complete data or audio unit.

Writing CD and DVD Data and Audio CDs

The following procedures are described in this section:

- “Restricting User Access to Removable Media With RBAC” on page 38
- “How to Restrict User Access to Removable Media With RBAC” on page 39
- “How to Identify a CD or DVD Writer” on page 39
- “How to Check the CD or DVD Media” on page 40
- “Creating a Data CD or DVD” on page 40
- “How to Create an ISO 9660 File System for a Data CD or DVD” on page 41
- “How to Create a Multi-Session Data CD” on page 41
- “Creating an Audio CD” on page 43
- “How to Create an Audio CD” on page 44
- “How to Extract an Audio Track on a CD” on page 45
- “How to Erase CD-RW Media” on page 46

The process of writing to a CD or DVD cannot be interrupted and needs a constant stream of data. Consider using the `cd rw -S` option to simulate writing to the media to verify that the system can provide data at a sufficient rate for writing to the CD or DVD.

Write errors can be caused by one of the following problems:

- The media cannot handle the drive speed. For example, some media are only certified for 2x or 4x speeds.
- The system is running too many heavy processes that are starving the writing process.
- The image is on a remote system, and network congestion is causing delays in reading the image.
- The source drive is slower than the destination drive.

If any of these problems occur, you can lower the writing speed of the device by using the `cdrw -p` option.

For example, the following command shows how to simulate writing at 4x speed:

```
$ cdrw -iS -p 4 image.iso
```

Note – CD-R, CD-RW (not MRW formatted), DVD-R, and DVD-RW media support simulation mode (-S), but DVD-RAM, DVD+R, DVD+RW, and any MRW-formatted media and some others do not support simulation mode. The following message is displayed if simulation mode is not supported:

```
Media does not support simulated writing
```

For more information about media types, see “[CD/DVD Media Commonly Used Terms](#)” on [page 36](#).

For more information, see [cdrw\(1\)](#).

Restricting User Access to Removable Media With RBAC

By default, all users can access removable media. However, you can restrict user access to removable media by setting up a role through role-based access control (RBAC). Access to removable media is restricted by assigning the role to a limited set of users.

For a discussion of using roles, see “[Role-Based Access Control \(Overview\)](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

▼ How to Restrict User Access to Removable Media With RBAC

1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Set up a role that includes the Device Management rights.

For more information, see Chapter 9, “Using Role-Based Access Control (Tasks),” in *Oracle Solaris 11.1 Administration: Security Services*.

```
# roleadd -m -d /export/home/muser -c "mediauser role" -A solaris.device.cdrw -P All muser
```

3 Add users who need to use the `cdrw` command to the newly created role.

```
# usermod -R muser joe
```

4 Comment the following line in the `/etc/security/policy.conf` file:

```
AUTHS_GRANTED=solaris.device.cdrw
```

If you do not do this step, all users still have access to the `cdrw` command, not just the members of the device management role.

After this file is modified, the device management role members are the only users who can use the `cdrw` command. Everyone else is denied access with the following message:

```
Authorization failed, Cannot access disks.
```

▼ How to Identify a CD or DVD Writer

1 Identify the CD or DVD writers on the system.

For example:

```
$ cdrw -l
Looking for CD devices...
  Node                | Connected Device                | Device type
-----+-----+-----
  cdrom0              | YAMAHA CRW8824S                | 1.0d | CD Reader/Writer
```

2 Identify a specific CD or DVD writer.

For example:

```
$ cdrw -a filename.wav -d cdrom2
```

3 Identify whether the media is blank or whether a table of contents exists on the media.

For example:

```
$ cdwr -M

Device : YAMAHA   CRW8824S
Firmware : Rev. 1.00 (26/04/00)
Media is blank
%
```

▼ How to Check the CD or DVD Media

The `cdwr` command works with or without removable media services running. For more information about disabling or enabling removable media services, see [“How to Disable or Enable Removable Media Services”](#) on page 28.

1 Insert a CD or DVD into the drive.

The CD or DVD can be any CD or DVD that the drive can read.

2 Check that the drive is connected properly by listing the drive.

```
$ cdwr -l
Looking for CD devices...
  Node                               Connected Device                               Device type
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
cdrom1                               | YAMAHA   CRW8824S                               1.0d | CD Reader/Writer
```

3 (Optional) If you do not see the drive in the list, select one of the following so that the system recognizes the drive.

- Add the drive without rebooting the system.

```
# devfsadm
```

Then restart removable media services.

Creating a Data CD or DVD

Prepare the data first by using the `mkisofs` command to convert the file and file information into the High Sierra format used on CDs or DVDs.

▼ How to Create an ISO 9660 File System for a Data CD or DVD

1 Insert a blank CD or DVD into the drive.

2 Create the ISO 9660 file system on the new CD or DVD.

```
$ mkisofs -r /pathname > cd-file-system
```

-r Creates Rock Ridge information and resets file ownerships to zero.

/pathname Identifies the path name used to create the ISO 9660 file system.

> cd-file-system Identifies the name of the file system to be put on the CD or DVD.

3 Copy the file system onto the CD or DVD.

```
$ cdrw -i cd-file-system
```

The -i *cd-file-system* specifies the image file for creating a data CD or DVD.

Example 2-1 Creating an ISO 9660 File System for a Data CD or DVD

The following example shows how to create an ISO 9660 file system for a data CD or DVD.

```
$ mkisofs -r /home/dubs/dir > dubs_cd
Total extents actually written = 56
Total translation table size: 0
Total rockridge attributes bytes: 329
Total directory bytes: 0
Path table size(bytes): 10
Max brk space used 8000
56 extents written (0 Mb)
```

Then, copy the file system onto the CD.

```
$ cdrw -i dubs_cd
Initializing device...done.
Writing track 1...done.
Finalizing (Can take several minutes)...done.
```

▼ How to Create a Multi-Session Data CD

This procedure describes how to put more than one session on a CD. This procedure includes an example of copying the `infoA` and `infoB` directories onto the CD.

1 Create the file system for the first CD session.

```
$ mkisofs -o infoA -r -V my_infoA /data/infoA
Total translation table size: 0
Total rockridge attributes bytes: 24507
Total directory bytes: 34816
Path table size(bytes): 98
Max brk space used 2e000
8929 extents written (17 Mb)
```

- o infoA Identifies the name of the ISO file system.
- r Creates Rock Ridge information and resets file ownerships to zero.
- V my_infoA Identifies a volume label to be used as the mount point by removable media services.
- /data/infoA Identifies the ISO image directory to create.

2 Copy the ISO file system for the first session onto the CD.

```
$ cdwr -i0 infoA
Initializing device...done.
Writing track 1...done.
done.
Finalizing (Can take several minutes)...done.
```

- i infoA Identifies the name of the image file to write to the CD.
- 0 Keeps the CD open for writing.

3 Reinsert the CD after it is ejected.

4 Identify the path name of the CD media to include in the next write session.

```
$ eject -n
.
.
.
cdrom0 -> /vol/dev/rdisk/c2t4d0/my_infoA
```

Note the /vol/dev/. . . path name.

5 Identify the next writeable address on the CD to write the next session.

```
% cdwr -M /cdrom
Device : YAMAHA CRW8424S
Firmware : Rev. 1.0d (06/10/99)

Track No. |Type |Start address
-----+-----+-----
1 |Audio |0
2 |Audio |33057
3 |Data |60887
4 |Data |68087
5 |Data |75287
Leadout |Data |84218
```

Last session start address: 75287
 Next writable address: 91118

Note the address in the Next writable address output so that you can provide this address when you write the next session.

6 Create the next ISO file system for the next CD session, and write it onto the CD.

```
$ mkisofs -o infoB -r -C 0,91118 -M /vol/dev/rdisk/c2t4d0/my_infoA
/data/infoB
Total translation table size: 0
Total rockridge attributes bytes: 16602
Total directory bytes: 22528
Path table size(bytes): 86
Max brk space used 20000
97196 extents written (189 Mb)
```

-o <i>infoB</i>	Identifies the name of the ISO file system.
-r	Creates Rock Ridge information and resets file ownerships to zero.
-C <i>0,91118</i>	Identifies the starting address of the first session and the next writable address.
-M <i>/vol/dev/rdisk/c2t4d0/my_infoA</i>	Specifies the path of the existing ISO image to be merged.
<i>/data/infoB</i>	Identifies the ISO image directory to create.

Creating an Audio CD

You can use the `cdwr` command to create audio CDs from individual audio tracks or from `.au` and `.wav` files.

The supported audio formats are described in the following table.

Format	Description
sun	Oracle <code>.au</code> file with data in Red Book CDDA format
wav	RIFF (<code>.wav</code>) file with data in Red Book CDDA format
cda	<code>.cda</code> file with raw CD audio data, which is 16-bit PCM stereo at 44.1 kHz sample rate in little-endian byte order
aur	<code>.aur</code> files with raw CD data in big-endian byte order

If no audio format is specified, the `cdwr` command tries to determine the audio file format based on the file extension. The case of the characters in the extension is ignored.

▼ How to Create an Audio CD

This procedure describes how to copy audio files onto a CD.

- 1 **Insert a blank CD into the CD-RW drive.**
- 2 **Change to the directory that contains the audio files.**

```
$ cd /myaudiodir
```

- 3 **Copy the audio files onto the CD.**

```
$ cdwr -a track1.wav track2.wav track3.wav
```

The `-a` option creates an audio CD.

Example 2-2 Creating an Audio CD

The following example shows how to create an audio CD.

```
$ cdwr -a bark.wav chirp.au meow.wav
Initializing device...done.
Writing track 1...done.
done.
Writing track 2...done.
Writing track 3...done.
done.
Finalizing (Can take several minutes)...done.
```

The following example shows how to create a multi-session audio CD. The CD is ejected after the first session is written. You would need to reinsert the CD before the next writing session.

```
$ cdwr -a0 groucho.wav chico.au harpo.wav
Initializing device...done.
Writing track 1...done.
done.
Writing track 2...done.
Writing track 3...done.
done.
Finalizing (Can take several minutes)...done.
<Re-insert CD>
$ cdwr -a zeppo.au
Initializing device...done.
Writing track 1...done.
done.
Finalizing (Can take several minutes)...done.
```

▼ How to Extract an Audio Track on a CD

Use the following procedure to extract an audio track from a CD and copy the audio track to a new CD.

If you don't use the `cdwr -T` option to specify the audio file type, the `cdwr` command uses the filename extension to determine the audio file type. For example, the `cdwr` command detects that this file is a `.wav` file.

```
$ cdwr -x 1 testme.wav
```

1 Insert an audio CD into the CD-RW drive.

2 Extract an audio track.

```
$ cdwr -x -T audio-type 1 audio-file
```

`-x` Extracts audio data from an audio CD.

`T audio-type` Identifies the type of audio file to be extracted. Supported audio types are `sun`, `wav`, `cda`, or `aur`.

`audio-file` Identifies the audio track to be extracted.

3 Copy the track to a new CD.

```
$ cdwr -a audio-file
```

Example 2-3 Extracting and Creating Audio CDs

The following example shows how to extract the first track from an audio CD and name the file `song1.wav`.

```
$ cdwr -x -T wav 1 song1.wav
Extracting audio from track 1...done.
```

This example shows how to copy a track to an audio CD.

```
$ cdwr -a song1.wav
Initializing device...done.
Writing track 1...done.
Finalizing (Can take several minutes)...done.
```

▼ How to Copy a CD

This procedure describes how to extract all the tracks from an audio CD into a directory and then copy all of them onto a blank CD.

Note – By default, the `cdwr` command copies the CD into the `/tmp` directory. The copying might require up to 700 MB of free space. If there is insufficient space in the `/tmp` directory for copying the CD, use the `-m` option to specify an alternate directory.

1 Insert an audio CD into a CD-RW drive.

2 Create a directory for the audio files.

```
$ mkdir /music_dir
```

3 Extract the tracks from the audio CD.

```
$ cdwr -c -m music_dir
```

An Extracting audio . . . message is display for each track.

The CD is ejected when all the tracks are extracted.

4 Insert a blank CD and press Return.

After the tracks are extracted, the audio CD is ejected. You are prompted to insert a blank CD.

Example 2-4 Copying a CD

This example shows how to copy one CD to another CD. You must have two CD-RW devices to do this task.

```
$ cdwr -c -s cdrom0 -d cdrom1
```

▼ How to Erase CD-RW Media

You have to erase existing CD-RW data before the CD can be rewritten.

● **Erase the entire media or just the last session on the CD by selecting one of the following:**

- Erase the last session only.

```
$ cdwr -d cdrom0 -b session
```

Erasing just the last session with the `-b session` option is faster than erasing the entire media with the `-b all` option. You can use the `-b session` option even if you used the `cdwr` command to create a data or audio CD in just one session.

- Erase the entire media.

```
$ cdwr -d cdrom0 -b all
```

Managing Devices (Tasks)

Device management in the Oracle Solaris release usually involves adding and removing peripheral devices from systems, possibly adding a third-party device driver to support a device, and displaying system configuration information. This chapter provides overview information and step-by-step instructions for managing peripheral devices, such as disks, DVD drives, and tape devices, in the Oracle Solaris release.

This is a list of the information in this chapter:

- “What's New in Device Management?” on page 47
- “Where to Find Additional Device Management Tasks” on page 48
- “Managing Devices in Oracle Solaris” on page 48
- “Adding a Peripheral Device to a System” on page 58
- “Accessing Devices” on page 60

What's New in Device Management?

This section provides information about new device management features in the Oracle Solaris release. For a complete listing of new Oracle Solaris features and a description of Oracle Solaris releases, see *Oracle Solaris 11.1 Release Notes*.

Support for USB 3.0 Devices

USB 3.0 support is provided in this Oracle Solaris release. A new USB host controller driver, `xhci`, is introduced. For more information, see [Chapter 5, “Managing USB Devices \(Tasks\)”](#).

Where to Find Additional Device Management Tasks

The following table describes where to find step-by-step instructions for hot-plugging devices and adding serial devices, such as printers and modems.

TABLE 3-1 Where to Find Instructions for Adding a Device

Device Management Task	For More Information
Add a disk that is not hot-pluggable.	Chapter 10, “Setting Up Disks (Tasks),” or “x86: Setting Up Disks for ZFS File Systems (Task Map)” on page 215
Hot-plug a SCSI or PCI device.	“SCSI Hot-Plugging With the <code>cfgadm</code> Command” on page 72 or “PCI or PCIe Hot-Plugging With the <code>cfgadm</code> Command” on page 82
Hot-plug a USB device.	“Managing USB Mass Storage Devices” on page 109
Add a modem.	Chapter 5, “Managing the System Console, Terminal Devices, and Power Services (Tasks),” in <i>Managing System Information, Processes, and Performance in Oracle Solaris 11.1</i>
Add a printer.	Chapter 1, “Setting Up and Administering Printers by Using CUPS (Overview),” in <i>Configuring and Managing Printing in Oracle Solaris 11.1</i>
Secure a device.	Chapter 5, “Controlling Access to Devices (Tasks),” in <i>Oracle Solaris 11.1 Administration: Security Services</i>

Managing Devices in Oracle Solaris

The following sections provide overview information about features that manage devices in Oracle Solaris:

- “About Device Drivers” on page 49
- “Automatic Configuration of Devices” on page 50
- “Displaying Device Configuration Information” on page 51
- “Resolving Faulty Devices” on page 55

For information about accessing devices, see “Accessing Devices” on page 60.

x86: Identifying Device Support

You can use the device detection tool to identify whether your x86 hardware is supported in this Oracle Solaris release. For more information, go to http://www.oracle.com/webfolder/technetwork/hcl/hcts/device_detect.jsp.

About Device Drivers

A computer typically uses a wide range of peripheral devices and mass-storage devices. Your system, for example, probably has a disk drive, a keyboard and a mouse, and some kind of DVD writer.

Other commonly used devices include the following:

- DVD drives
- Printers and plotters
- Light pens
- Touch-sensitive screens
- Digitizers
- Tablet-and-stylus pairs

The Oracle Solaris software does not directly communicate with all of these devices. Each type of device requires different data formats, protocols, and transmission rates.

A *device driver* is a low-level program that enables the operating system to communicate with a specific piece of hardware. The driver serves as the operating system's “interpreter” for that piece of hardware.

▼ How to Customize a Driver Configuration

In the Oracle Solaris 11 release, driver customizations are made in the `/etc/driver/drv` directory rather than in the `/kernel` directory as in previous releases. This improvement means that your driver customizations are not overwritten when the system is upgraded. The files in the `/etc/driver/drv` directory are preserved during the upgrade.

Customizing a driver configuration usually means that a per-device parameter or global property that impacts all devices is added or modified.

- 1 **Become an administrator.**
- 2 **Copy the original vendor supplied *driver.conf* file to the `/etc/driver/drv` directory. For example:**

```
# cp /kernel/drv/sd.conf /etc/driver/drv/sd.conf
```

- 3 **Modify the parameter entry and save the file.**

For example, the `sd.conf` includes the following entry for `sd` device at target 0, lun 0:

```
name="sd" class="scsi" target=0 lun=0;
```

To add the `retries` parameter for this device, modify the existing entry as follows:

```
name="sd" class="scsi" target=0 lun=0 retries=4;
```

4 Display the customized property value. For example:

```
# prtconf -vu
sd, instance #1
    Admin properties:
        name='retries' type=int items=1
        value=00000004
```

Automatic Configuration of Devices

The kernel consists of a small generic core with a platform-specific component and a set of modules. The kernel is configured automatically in the Oracle Solaris release.

A *kernel module* is a software component that is used to perform a specific task on the system. An example of a *loadable* kernel module is a device driver that is loaded when the device is accessed.

The kernel modules are described in the following table.

TABLE 3-2 Description of Solaris Kernel Modules

Location	Directory Contents
/platform/'uname -m'/kernel	Platform-specific kernel components
/kernel	Kernel components common to all platforms that are needed for booting the system
/usr/kernel	Kernel components common to all platforms within a particular instruction set

The system determines what devices are attached to it at boot time. Then, the kernel configures itself dynamically, loading needed modules into memory. At this time, device drivers are loaded when devices, such as disk devices and tape devices, are accessed. This process is called *autoconfiguration* because all kernel modules are loaded automatically when they are needed.

You can customize the way in which kernel modules are loaded by modifying the `/etc/system` file. For instructions on modifying this file, see [system\(4\)](#).

Features and Benefits of Autoconfiguration

The benefits of autoconfiguration are as follows:

- Main memory is used more efficiently because modules are loaded when needed.
- There is no need to reconfigure the kernel when new devices are added to the system.
- Drivers can be loaded and tested without having to rebuild the kernel.

Autoconfiguration is used when you add a new device (and driver) to the system. In previous releases, it was necessary to perform a reconfiguration boot if you added a device to a system that is shutdown. Device configuration enhancements make a reconfiguration boot unnecessary when a device is added to a system that is shutdown.

You can add, remove, or replace devices in the Oracle Solaris OS while the system is still running, if the system components support hot-plugging. For information about hot-plugging devices, see [Chapter 4, “Dynamically Configuring Devices \(Tasks\)”](#).

What You Need for Unsupported Devices

Device drivers that are needed to support a wide range of standard devices are included in the Oracle Solaris release. These drivers can be found in the `/kernel/drv` and `/platform/uname-m/kernel/drv` directories.

However, if you have purchased an unsupported device, the manufacturer should provide the software that is needed for the device to be properly installed, maintained, and administered.

At a minimum, this software includes a device driver and its associated configuration (`.conf`) file. The `.conf` files reside in the `drv` directories. This software might also include custom maintenance and administrative utilities because the device might be incompatible with Oracle Solaris utilities.

For more information about what you need for unsupported devices, contact your device manufacturer.

Displaying Device Configuration Information

The following table describes the three commands that are used to display system and device configuration information.

Command	Description	Man Page
<code>prtconf</code>	Displays system configuration information, including the total amount of memory and the device configuration, as described by the system's device hierarchy. The output that is displayed by this command depends upon the type of system.	prtconf(1M)
<code>sysdef</code>	Displays device configuration information, including system hardware, pseudo devices, loadable modules, and selected kernel parameters.	sysdef(1M)
<code>dmesg</code>	Displays system diagnostic messages, as well as a list of devices that are attached to the system since the last reboot.	dmesg(1M)

For information about the device names that are used to identify devices on the system, see [“Device Naming Conventions”](#) on page 61.

driver not attached Message

The following driver-related message might be displayed by the `prtconf` and `sysdef` commands:

```
device, instance #number (driver not attached)
```

This message does not always mean that a driver is unavailable for this device. This message means that no driver is *currently* attached to the device instance because no device exists at this node or the device is not in use. Drivers are loaded automatically when the device is accessed. They are unloaded when the device is not in use.

In-Use Device Error Checking

The following utilities detect when a specified device is in use:

- `dumpadm`
- `format`
- `mkfs` and `newfs`
- `swap`

These enhancements mean that these utilities might detect some of the following usage scenarios:

- Device is part of a ZFS storage pool
- Device is a dump or swap device
- Mounted file system or an entry for the device exists in the `/etc/vfstab` file

For example, if you attempt to use the `format` utility to access an active device, you will see a message similar to the following:

```
# format
Searching for disks..done
AVAILABLE DISK SELECTIONS:
  0. clt0d0 <FUJITSU-MAY2073RCSUN72G-0401 cyl 8921 alt 2 hd 255 sec 63>
    /pci@0,0/pci1022,7450@2/pci1000,3060@3/sd@0,0
    /dev/chassis/SYS/HD0/disk
  1. clt1d0 <FUJITSU-MAY2073RCSUN72G-0401-68.37GB>
    /pci@0,0/pci1022,7450@2/pci1000,3060@3/sd@1,0
    /dev/chassis/SYS/HD1/disk
Specify disk (enter its number): 0
selecting clt0d0
[disk formatted]
/dev/dsk/clt0d0s0 is part of active ZFS pool rpool. Please see zpool(1M).
```

```
FORMAT MENU:
```

```
.
.
.
```

▼ How to Display System Configuration Information

Use the output of the `prtconf` and `sysdef` commands to identify which disk, tape, and DVD devices are connected to the system. The output of these commands displays the driver not attached messages next to the device instances. Because these devices are always being monitored by some system process, the driver not attached message is usually a good indication that no device exists at that device instance.

Use the `sysdef` command to display system configuration information that includes pseudo devices, loadable modules, and selected kernel parameters.

● Display system and device configuration information.

■ Display all the devices connected to a system.

For example, the following `prtconf -v` output on a SPARC system identifies the disk devices connected to the system. The detailed disk information is described in the Device Minor Nodes section.

```
$ /usr/sbin/prtconf -v | more
Memory size: 32640 Megabytes
System Peripherals (Software Nodes):

SUNW,SPARC-Enterprise-T5220
.
.
.
Location: /dev/chassis/SUN-Storage-J4400.0918QAKA24/SCSI_Device__2/disk
Device Minor Nodes:
  dev=(27,40)
    dev_path=/pci@0/pci@0/pci@8/pci@0/pci@a/LSILogic,sas@0/sd@2,0:a
    spectype=blk type=minor
    dev_link=/dev/dsk/c4t2d0s0
  dev_path=/pci@0/pci@0/pci@8/pci@0/pci@a/LSILogic,sas@0/sd@2,0:a,raw
    spectype=chr type=minor
    dev_link=/dev/rdisk/c4t2d0s0
  Device Minor Layered Under:
    mod=zfs accesstype=blk
    dev_path=/pseudo/zfs@0
  Minor properties:
    name='Nblocks' type=int64 items=1 dev=(27,40)
      value=0000000074702c8f
    name='Size' type=int64 items=1 dev=(27,40)
      value=000000e8e0591e00
.
.
.
```

■ Display information about one specific device connected to the system.

For example, the following `prtconf` output on a SPARC system displays the `sd` instance number for `/dev/dsk/c4t2d0s0`.

```
# prtconf -v /dev/dsk/c4t2d0s0 | grep instance
sd, instance #5
```

- **Display only the devices that are attached to the system.**

```
# prtconf | grep -v not
```

- **Display device usage information.**

For example, the following `fuser` command displays which processes are accessing the `/dev/console` device.

```
# fuser -d /dev/console
/dev/console:  5742o   2269o   20322o   858o
```

Example 3-1 Displaying System Configuration Information

The following `prtconf` output is displayed on an x86 based system.

```
# prtconf
System Configuration: Oracle Corporation i86pc
Memory size: 8192 Megabytes
System Peripherals (Software Nodes):

i86pc
  scsi_vhci, instance #0
  pci, instance #0
    pci108e,4843, instance #0
    pci8086,25e2, instance #0
      pci8086,3500, instance #7
        pci8086,3510, instance #9
        pci8086,3518, instance #10
          pci108e,4843, instance #0
          pci108e,4843, instance #1
        pci8086,350c, instance #8
      pci8086,25e3 (driver not attached)
      pci8086,25f8, instance #2
        pci108e,286, instance #0
          disk, instance #0
          disk, instance #2
          disk, instance #3
          disk, instance #1
        pci8086,25e5 (driver not attached)
        pci8086,25f9 (driver not attached)
        pci8086,25e7 (driver not attached)
        pci108e,4843, instance #0 (driver not attached)
        pci108e,4843, instance #1
        pci108e,4843, instance #2 (driver not attached)
        pci108e,4843 (driver not attached)
        pci108e,4843 (driver not attached)
        pci108e,4843 (driver not attached)
        pci108e,4843 (driver not attached)
        pci8086,2690, instance #6
          pci108e,125e, instance #2
          pci108e,125e, instance #3
        pci108e,4843, instance #0
        pci108e,4843, instance #1
          device, instance #0
            keyboard, instance #0
            mouse, instance #1
```

```

        pci108e,4843, instance #2
        pci108e,4843, instance #3
        pci108e,4843, instance #0
            storage, instance #0
            disk, instance #4
    .
    .
    .

```

The following sysdef output is displayed from an x86 based system:

```

# sysdef
* Hostid
*
    29f10b4d
*
* i86pc Configuration
*
*
* Devices
*
+boot (driver not attached)
memory (driver not attached)
aliases (driver not attached)
chosen (driver not attached)
i86pc-memory (driver not attached)
i86pc-mmio (driver not attached)
openprom (driver not attached)
options, instance #0
packages (driver not attached)
delayed-writes (driver not attached)
itu-props (driver not attached)
isa, instance #0
    motherboard (driver not attached)
    pnpADP,1542, instance #0
    asy, instance #0
    asy, instance #1
    lp, instance #0 (driver not attached)
    fd, instance #0
        fd, instance #0
        fd, instance #1 (driver not attached)
    kd (driver not attached)
    kdmouse (driver not attached)
.
.
.

```

Resolving Faulty Devices

A device retirement mechanism isolates a device as *faulty* by the fault management framework (FMA). This feature allows faulty devices to be safely and automatically inactivated to avoid data loss, data corruption, or panics and system down time. The retirement process is done safely, taking into account the stability of the system after the device has been retired.

Critical devices are never retired. If you need to manually replace a retired device, use the `fmadm repair` command after the device replacement so that system knows that the device is replaced, in addition to the manual replacement steps.

For more information, see [fmadm\(1M\)](#).

A general message regarding device retirement is displayed on the console and written to the `/var/adm/messages` file so that you are aware of a retired device. For example:

```
Aug 9 18:14 starbug genunix: [ID 751201 kern.notice] NOTICE: One or more I/O devices have been retired
```

You can use the `prtconf` command to identify specific retired devices. For example:

```
# prtconf
.
.
.
pci, instance #2
    scsi, instance #0
        disk (driver not attached)
        tape (driver not attached)
        sd, instance #3
        sd, instance #0 (retired)
    scsi, instance #1 (retired)
        disk (retired)
        tape (retired)
pci, instance #3
    network, instance #2 (driver not attached)
    network, instance #3 (driver not attached)
os-io (driver not attached)
iscsi, instance #0
pseudo, instance #0
.
.
.
```

▼ How to Resolve a Faulty Device

Use the steps that follow to resolve a faulty device or a device that has been retired.

You can also review ZFS device problem or failure information by using the `zpool status` or the `fmadm` command. For ZFS device problem or failure information, see [Chapter 10, “Oracle Solaris ZFS Troubleshooting and Pool Recovery,”](#) in *Oracle Solaris 11.1 Administration: ZFS File Systems*.

1 Identify the faulted device with the `fmadm faulty` command. For example:

```
# fmadm faulty
-----
TIME          EVENT-ID          MSG-ID SEVERITY
-----
Jun 20 16:30:52 55c82fff-b709-62f5-b66e-b4e1bbe9dcb1 ZFS-8000-LR Major
```

```

Problem Status      : solved
Diag Engine        : zfs-diagnosis / 1.0
System
  Manufacturer     : unknown
  Name             : ORCL,SPARC-T3-4
  Part_Number      : unknown
  Serial_Number    : 1120BDRCCD
  Host_ID          : 84a02d28

```

```

-----
Suspect 1 of 1 :
  Fault class : fault.fs.zfs.open_failed
  Certainty  : 100%
  Affects    : zfs://pool=86124fa573cad84e/vdev=25d36cd46e0a7f49/
pool_name=pond/vdev_name=id1,sd@n5000c500335dc60f/a
  Status     : faulted and taken out of service

FRU
  Name       : "zfs://pool=86124fa573cad84e/vdev=25d36cd46e0a7f49/
pool_name=pond/vdev_name=id1,sd@n5000c500335dc60f/a"
  Status    : faulty

Description : ZFS device 'id1,sd@n5000c500335dc60f/a' in pool 'pond' failed to
              open.

Response    : An attempt will be made to activate a hot spare if available.

Impact     : Fault tolerance of the pool may be compromised.

Action     : Use 'fmadm faulty' to provide a more detailed view of this event.
              Run 'zpool status -lx' for more information. Please refer to the
              associated reference document at
              http://support.oracle.com/msg/ZFS-8000-LR for the latest service
              procedures and policies regarding this diagnosis.

```

2 Replace the faulty or retired device or clear the device error. For example:

```
# zpool clear pond c0t5000C500335DC60Fd0
```

If an intermittent device error occurred but the device was not replaced, you can attempt to clear the previous error.

3 Clear the FMA fault. For example:

```
# fmadm repaired zfs://pool=86124fa573cad84e/vdev=25d36cd46e0a7f49/ \
pool_name=pond/vdev_name=id1,sd@n5000c500335dc60f/a
fmadm: recorded repair to of zfs://pool=86124fa573cad84e/vdev=25d36cd46e0a7f49/
pool_name=pond/vdev_name=id1,sd@n5000c500335dc60f/a
```

4 Confirm that the fault is cleared.

```
# fmadm faulty
```

If the error is cleared, the `fmadm faulty` command returns nothing.

Adding a Peripheral Device to a System

Adding a new peripheral device that is not hot-pluggable usually involves the following:

- Shutting down the system
- Connecting the device to the system
- Rebooting the system

Use “[How to Add a Peripheral Device](#)” on page 58 to add the following devices that are not hot-pluggable to a system:

- DVD drive
- Secondary disk drive
- Tape drive

In some cases, you might have to add a third-party device driver to support the new device.

For information on hot-plugging devices, see [Chapter 4, “Dynamically Configuring Devices \(Tasks\)”](#).

▼ How to Add a Peripheral Device

1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

2 (Optional) If you need to add a device driver to support the device, complete the procedure “[How to Add a Device Driver](#)” on page 59.

3 Shut down the system.

```
# shutdown -i0 -g30 -y
```

-i0 Brings the system to the 0 init state, which is the appropriate state for turning the system power off for adding and removing devices.

-g30 Shuts the system down in 30 seconds. The default is 60 seconds.

-y Continues the system shutdown without user intervention. Otherwise, you are prompted to continue the shutdown process.

4 Select one of the following to turn off power to the system after it is shut down:

- For SPARC platforms, it is safe to turn off power if the ok prompt is displayed.
- For x86 platforms, it is safe to turn off power if the type any key to continue prompt is displayed.

5 Turn off power to all peripheral devices.

For the location of power switches on any peripheral devices, refer to the hardware installation guides that accompany your peripheral devices.

6 Install the peripheral device, making sure that the hardware configuration adheres to manufacturer specifications.

Refer to the hardware installation guide that accompanies the peripheral device for information on installing and connecting the device.

7 Turn on the power to the system.

The system boots to multiuser mode, and the login prompt is displayed.

8 Verify that the peripheral device has been added by attempting to access the device.

For information on accessing the device, see [“Accessing Devices” on page 60](#).

▼ How to Add a Device Driver

In the Oracle Solaris 11 release, the `pkg` commands are used to add packages to the system. Device driver packages might still be in SRV4 package format so the steps below use the `pkgadd` command instead of the `pkg install` command.

This procedure assumes that the device has already been added to the system. If not, see [“How to Add a Peripheral Device” on page 58](#).

1 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*](#).

2 Place the media into the drive.**3 Install the driver.**

```
# pkgadd [-d] device package-name
```

4 Verify that the package has been added correctly.

```
# pkgchk package-name
#
```

The system prompt returns with no response if the package is installed correctly.

Example 3–2 Adding a Device Driver

The following example shows how to install and verify a package called XYZdrv.

```
# pkgadd XYZdrv
(licensing messages displayed)
.
.
.
Installing XYZ Company driver as <XYZdrv>
.
.
.
Installation of <XYZdrv> was successful.
# pkgchk XYZdrv
#
```

Accessing Devices

You need to know how to specify device names when using commands to manage disks, file systems, and other devices. In most cases, you can use logical device names to represent devices that are connected to the system. Both logical and physical device names are represented on the system by logical and physical device files.

How Device Information Is Created

When a system is booted for the first time, a device hierarchy is created to represent all the devices connected to the system. The kernel uses the device hierarchy information to associate drivers with their appropriate devices. The kernel also provides a set of pointers to the drivers that perform specific operations.

How Devices Are Managed

The `devfs` file system manages the `/devices` directory, which is the name space of all devices on the system. This directory represents the *physical* devices that consists of actual bus and device addresses.

The `dev` file system manages the `/dev` directory, which is the name space of *logical* device names.

By default, the `devfsadm` command attempts to load every driver in the system and attach to all possible device instances. Then, `devfsadm` creates the device files in the `/devices` directory and the logical links in the `/dev` directory. The `devfsadm` command also maintains the `path_to_inst` instance database.

Updates to the `/dev` and `/devices` directories in response to dynamic reconfiguration events or file system accesses are handled by `devfsd`, the daemon version of the `devfsadm` command. This daemon is started by the service management facility when a system is booted.

Because the `devfsadm` daemon automatically detects device configuration changes generated by any reconfiguration event, there is no need to run this command interactively.

For more information, see the following references:

- [devfsadm\(1M\)](#)
- [dev\(7FS\)](#)
- [devfs\(7FS\)](#)
- [path_to_inst\(4\)](#)

Device Naming Conventions

Devices are referenced in one of three ways in Oracle Solaris:

- **Physical device name** – Represents the full device path name in the device information hierarchy. The physical device name is created by when the device is first added to the system. Physical device files are found in the `/devices` directory.
- **Instance name** – Represents the kernel's abbreviation name for every possible device on the system. For example, `sd0` and `sd1` represent the instance names of two disk devices. Instance names are mapped in the `/etc/path_to_inst` file.
- **Logical device name** – The logical device name is created by when the device is first added to the system. Logical device names are used with most file system commands to refer to devices. For a list of file commands that use logical device names, see [Table 3–3](#). Logical device files in the `/dev` directory are symbolically linked to physical device files in the `/devices` directory.

The preceding device name information is displayed with the following commands:

- `dmesg`
- `format`
- `sysdef`
- `prtconf`

Logical Disk Device Names

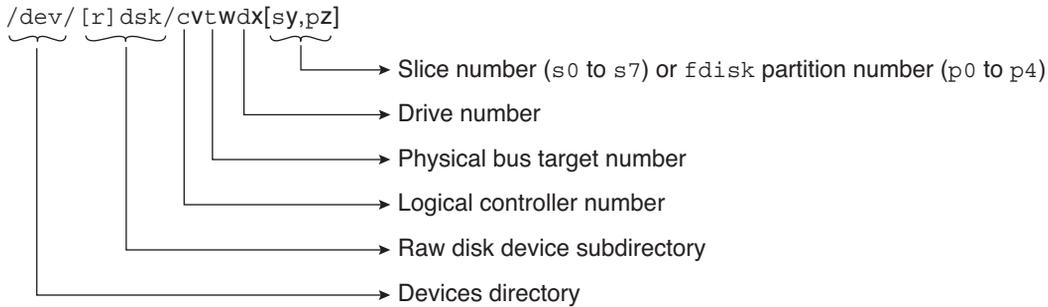
Logical device names are used to access disk devices when you perform the following tasks:

- Add a new disk to the system.
- Move a disk from one system to another system.
- Access or mount a file system residing on a local disk.
- Back up a local file system.

Many administration commands take arguments that refer to a disk slice or file system.

Refer to a disk device by specifying the subdirectory to which it is symbolically linked, either `/dev/dsk` or `/dev/rdisk`, followed by a string identifying the particular controller, disk, and slice.

FIGURE 3-1 Description of Logical Device Names



Specifying the Disk Subdirectory

Some disk and file administration commands require the use of either a *raw* (or *character*) device interface, or a *block* device interface. The distinction is made by how data is read from the device.

Raw device interfaces transfer only small amounts of data at a time. Block device interfaces include a buffer from which large blocks of data are read at once.

Different commands require different interfaces:

- When a command requires the raw device interface, specify the `/dev/rdsk` subdirectory. (The “r” in `rdsk` stands for “raw.”)
- When a command requires the block device interface, specify the `/dev/dsk` subdirectory.
- When you are not sure whether a command requires use of `/dev/dsk` or `/dev/rdsk`, check the man page for that command.

The following table shows which interface is required for some commonly used disk and file system commands.

TABLE 3-3 Device Interface Type Required by Some Frequently Used Commands

Command Reference	Interface Type	Example of Use
<code>dumpadm(1M)</code>	Block	<code>dumpadm -d /dev/zvol/dsk/rpool/dump</code>
<code>prtvtoc(1M)</code>	Raw	<code>prtvtoc /dev/rdsk/c0t0d0s0</code>

TABLE 3-3 Device Interface Type Required by Some Frequently Used Commands (Continued)

Command Reference	Interface Type	Example of Use
<code>swap(1M)</code>	Block	<code>swap -a /dev/zvol/dsk/rpool/swap</code>

Direct and Bus-Oriented Controllers

You might access disk partitions or slices differently depending upon whether the disk device is connected to a direct or bus-oriented controller. Generally, direct controllers do not include a *target* identifier in the logical device name.

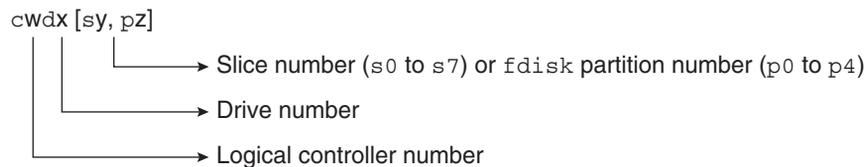
The conventions for both types of controllers are explained in the following subsections.

Note – Controller numbers are assigned automatically during system initialization. The numbers are strictly logical and imply no direct mapping to physical controllers.

Disks With Direct Controllers

To specify a slice on a disk with an IDE controller, follow the naming convention that is shown in the following figure.

FIGURE 3-2 Disks With Direct Controllers

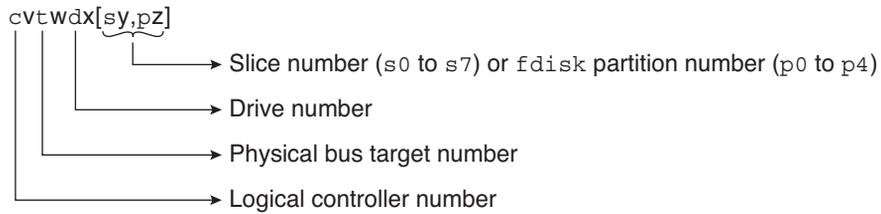


To indicate the entire `fdisk` partition, specify slice 2 (`s2`).

Disks With Bus-Oriented Controllers

To specify a slice on a disk with a bus-oriented controller, SCSI for instance, follow the naming convention that is shown in the following figure.

FIGURE 3-3 Disks With Bus-Oriented Controllers

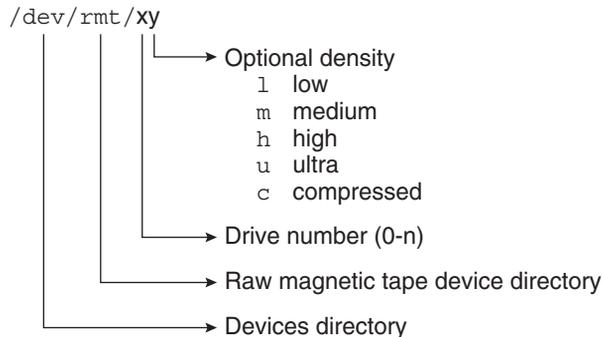


To indicate the whole disk, specify slice 2 (`s2`).

Logical Tape Device Names

Logical tape device files are found in the `/dev/rmt/*` directory as symbolic links from the `/devices` directory.

FIGURE 3-4 Logical Tape Device Names



The first tape device connected to the system is 0 (`/dev/rmt/0`). Tape density values (`l`, `m`, `h`, `c`, and `u`) are described in [Chapter 18, “Managing Tape Drives \(Tasks\)”](#).

Logical Removable Media Device Names

Since removable media is managed by removable media management services, the logical device name is usually not used unless you want to mount the media manually.

The logical device name that represents the removable media devices on a system are described in [“Accessing Removable Media” on page 27](#).

Dynamically Configuring Devices (Tasks)

This chapter provides instructions for dynamically configuring devices in the Oracle Solaris OS. You can add, remove, or replace devices in the Oracle Solaris OS while the system is still running, if the system components support hot-plugging. If the system components do not support hot-plugging, new devices will be configured at boot time, after the new components are installed in the system.

For information about the procedures that are associated with dynamically configuring devices, see the following sections:

- “SCSI Hot-Plugging With the `cfgadm` Command (Task Map)” on page 71
- “PCI or PCIe Hot-Plugging With the `cfgadm` Command (Task Map)” on page 81
- “SATA Hot-Plugging With the `cfgadm` Command” on page 88
- “Application Developer RCM Script (Task Map)” on page 91
- “System Administrator RCM Script (Task Map)” on page 92

For information on hot-plugging USB devices with the `cfgadm` command, see “Hot-Plugging USB Devices With the `cfgadm` Command” on page 129.

For information on hot-plugging InfiniBand devices with `cfgadm` command, see [Chapter 6](#), “Using InfiniBand Devices (Overview/Tasks).”

For information about accessing devices, see “Accessing Devices” on page 60.

Dynamic Reconfiguration and Hot-Plugging

The following information is described in this section:

- “Attachment Points” on page 67
- “Detaching PCI or PCIe Adapter Cards” on page 69
- “Attaching PCI or PCIe Adapter Cards” on page 69
- “PCIe Hot-Plugging With the (`hotplug`) Command” on page 69

Hot-plugging is the ability to physically add, remove, or replace system components while the system is running. *Dynamic reconfiguration* refers to the ability to hot-plug system components. This term also refers to the general ability to move system resources (both hardware and software) around in the system or to disable them in some way without physically removing them from the system.

Generally, you can hot-plug the following bus types:

- USB
- Fibre Channel
- 1394
- ATA
- SCSI
- Infiniband

In addition, you can hot-plug the following devices with the `cfgadm` command on both SPARC and x86 platforms:

- USB devices
- SCSI or SAS devices
- PCI devices
- PCIe devices
- SATA devices
- InfiniBand devices

Features of the `cfgadm` command include the following:

- Displaying system component status
- Testing system components
- Changing component configurations
- Displaying configuration help messages

The benefit of using the `cfgadm` command to reconfigure systems components is that you can add, remove, or replace components while the system is running. An added benefit is that the `cfgadm` command guides you through the steps needed to add, remove, or replace system components.

For step-by-step instructions on hot-plugging components, see the following:

- [“SCSI Hot-Plugging With the `cfgadm` Command” on page 72](#)
- [“PCI or PCIe Hot-Plugging With the `cfgadm` Command” on page 82](#)
- [“SATA Hot-Plugging With the `cfgadm` Command” on page 88](#)
- `cfgadm(1M)`

Note – Not all SCSI and PCI controllers support hot-plugging with the `cfgadm` command.

As part of Oracle's high availability strategy, dynamic reconfiguration is expected to be used in conjunction with additional layered products, such as alternate pathing or fail over software. Both products provide fault tolerance in the event of a device failure.

Without any high availability software, you can replace a failed device by manually stopping the appropriate applications, unmounting noncritical file systems, and then proceeding with the add or remove operations.

Note – Some systems have slots that hot-pluggable and slots that are not hot-pluggable. For information about hot-plugging devices on your specific hardware configuration, such as on enterprise-level systems, refer to your hardware configuration documentation.

Attachment Points

The `cfgadm` command displays information about *attachment points*, which are locations in the system where dynamic reconfiguration operations can occur.

An attachment point consists of the following:

- An *occupant*, which represents a hardware component that can be configured into the system
- A *receptacle*, which is the location that accepts the occupant

Attachment points are represented by logical and physical attachment point IDs (*Ap_Ids*). The physical *Ap_Id* is the physical path name of the attachment point. The logical *Ap_Id* is a user-friendly alternative for the physical *Ap_Id*. For more information on *Ap_Ids*, refer to [cfgadm\(1M\)](#).

The logical *Ap_Id* for a SCSI Host Bus Adapter (HBA), or SCSI controller, is usually represented by the controller number, such as `c0`.

In cases where no controller number has been assigned to a SCSI HBA, then an internally generated unique identifier is provided. An example of a unique identifier for a SCSI controller is the following:

```
fas1:scsi
```

The logical *Ap_Id* for a SCSI device usually has this format:

```
HBA-logical-apid::device-identifier
```

In the following example, `c0` is the logical *Ap_Id* for the SCSI HBA:

```
c0::dsk/c0t3d0
```

The device identifier is typically derived from the logical device name for the device in the /dev directory. For example, a tape device with logical device name, /dev/rmt/1, has the following logical Ap_Id:

```
c0: : rmt/1
```

If a logical Ap_Id of a SCSI device cannot be derived from the logical name in the /dev directory, then an internally generated unique identifier is provided. An example of an identifier for the /dev/rmt/1 tape device is the following:

```
c0: : st4
```

For more information on SCSI Ap_Ids, refer to [cfgadm_scsi\(1M\)](#).

The `cfgadm` command represents all resources and dynamic reconfiguration operations in terms of a common set of states (such as configured and unconfigured) and operations (such as connect, configure, unconfigure, and so on). For more information on these common states and operations, see [cfgadm\(1M\)](#).

The following table shows the receptacle and occupant states for the SCSI HBA attachment points.

Receptacle State	Description	Occupant State	Description
empty	N/A for SCSI HBA	configured	One or more devices is configured on the bus
disconnected	Bus quiesced	unconfigured	No devices are configured
connected	Bus active		

The following table shows the receptacle and occupant states for SCSI device attachment points.

Receptacle State	Description	Occupant State	Description
empty	N/A for SCSI devices	configured	Device is configured
disconnected	Bus quiesced	unconfigured	Device is not configured
connected	Bus active		

The state of SCSI attachment points is unknown unless special hardware indicates otherwise. For instructions on displaying SCSI component information, see [“How to Display Information About SCSI Devices”](#) on page 72.

Detaching PCI or PCIe Adapter Cards

A PCI adapter card that is hosting nonvital system resources can be removed if the device driver supports hot-plugging. A PCI adapter card is not detachable if it is a vital system resource.

For a PCI adapter card to be detachable, the following conditions must be met:

- The device driver must support hot-plugging.
- Critical resources must be accessible through an alternate pathway.

For example, if a system has only one Ethernet card installed in it, the Ethernet card cannot be detached without losing the network connection. This detachment requires additional layered software support to keep the network connection active.

Attaching PCI or PCIe Adapter Cards

A PCI adapter card can be added to the system as long as the following conditions are met:

- There are slots available.
- The device driver supports hot-plugging for this adapter card.

For step-by-step instructions on adding or removing a PCI adapter card, see [“PCI or PCIe Hot-Plugging With the `cfgadm` Command”](#) on page 82.

PCIe Hot-Plugging With the (`hotplug`) Command

You can use the `hotplug` command to manage hot pluggable connections, where a connection can be a connector or port, on PCI Express (PCIe) and PCI SHPC devices only. A *hotplug connector* is a physical representation in the system where a component is inserted or removed. A *hotplug port* is a logical representation in the system device tree where the connection of a device to the system is managed.

You can use the `hotplug` features to online or offline a device, even an on-board device, without physically adding or removing the device from the system.

You must enable the `hotplug` service to manage devices with the `hotplug` command.

```
# svcadm enable svc:/system/hotplug:default
```

The following examples show how to use the `hotplug` command:

Display all the PCI/PCIe hot-pluggable connectors/ports (virtual and physical) in the system as follows:

```
# hotplug list -lv
```

Configure an Ethernet card in a PCIe slot. For example:

```
# hotplug enable /pci0,0 pcie0
```

Unconfigure an Ethernet card in a PCIe slot. For example:

```
# hotplug disable /pci0,0 pcie0
```

Offline a PCI device node, which means detach the device driver for that node. For example:

```
# hotplug offline /pci0,0/pci1 pci.0,2
```

Online a PCI device node, which means attach the driver for that node. For example:

```
# hotplug online /pci0,0/pci1 pci.0,2
```

Install dependent ports of an IOV physical function. For example:

```
# hotplug install /pci@400/pci@1/pci@0/pci@4 pci.0,1
```

Then, display the resulting IOV virtual functions that were probed. For example:

```
# hotplug list -v /pci@400/pci@1/pci@0/pci@4 pci.0,1
<pci.0,1> (ONLINE)
  { IOV physical function }
  { IOV virtual function 'pci.0,81' }
  { IOV virtual function 'pci.0,83' }
  { IOV virtual function 'pci.0,85' }
  { IOV virtual function 'pci.0,87' }
<pci.0,81> (OFFLINE)
ethernet@0,81
<pci.0,83> (OFFLINE)
ethernet@0,83
<pci.0,85> (OFFLINE)
ethernet@0,85
<pci.0,87> (OFFLINE)
ethernet@0,87
```

Uninstall the dependent ports of an IOV physical function. For example:

```
# hotplug uninstall /pci@400/pci@1/pci@0/pci@4 pci.0,0
```

This operation fails if a dependent IOV virtual function is busy. For example:

```
# hotplug uninstall /pci@400/pci@1/pci@0/pci@4 pci.0,0
ERROR: devices or resources are busy.
ethernet@0,81:
  { Network interface igbvf1 }
  { igbvf1: hosts IP addresses: 10.0.0.1 }
  { Plumbed IP Address }
```

Troubleshooting PCI Hot Plug Operations (hotplug)

You might see the following maintenance states for an attached device in a hot-pluggable port.

```

/pci@0,0/pci10de,5d@e <pci.a,1> (MAINTENANCE)
.
.
./pci@0,0/pci108e,534a@d <pci.0,0> (MAINTENANCE-SUSPENDED)

```

These messages indicate that a fault event or a maintenance operation occurred. The MAINTENANCE states mean that a device is in use, but it is not fully operational. The MAINTENANCE-SUSPENDED state means that the device is *live suspended*, due to a maintenance operation. For example, reconfiguring the device hardware.

The following service must be running to use the `hotplug` command.

```
svc:/system/hotplug:default
```

Otherwise, you will see the following message:

```
ERROR: hotplug service is not available.
```

The following error message is displayed on systems that do not have any supported I/O buses:

```

ERROR: there are no connections to display.
      (See hotplug(1m) for more information.)

```

The above message could mean that the system might have other hot-pluggable I/O devices, but you need to use the `cfgadm` command rather than the `hotplug` command to manage these devices.

SCSI Hot-Plugging With the `cfgadm` Command (Task Map)

Task	Description	For Instructions
Display information about SCSI devices.	Display information about SCSI controllers and devices.	“How to Display Information About SCSI Devices” on page 72
Unconfigure a SCSI controller.	Unconfigure a SCSI controller.	“How to Unconfigure a SCSI Controller” on page 73
Configure a SCSI controller.	Configure a SCSI controller that was previously unconfigured.	“How to Configure a SCSI Controller” on page 74
Configure a SCSI device.	Configure a specific SCSI device.	“How to Configure a SCSI Device” on page 74
Disconnect a SCSI controller.	Disconnect a specific SCSI controller.	“How to Disconnect a SCSI Controller” on page 75
Connect a SCSI controller.	Connect a specific SCSI controller that was previously disconnected.	“SPARC: How to Connect a SCSI Controller” on page 76

Task	Description	For Instructions
Add a SCSI device to a SCSI bus.	Add a specific SCSI device to a SCSI bus.	“SPARC: How to Add a SCSI Device to a SCSI Bus” on page 76
Replace an identical device on a SCSI controller.	Replace a device on the SCSI bus with another device of the same type.	“SPARC: How to Replace an Identical Device on a SCSI Controller” on page 77
Remove a SCSI device.	Remove a SCSI device from the system.	“SPARC: How to Remove a SCSI Device” on page 79
Troubleshoot SCSI configuration problems.	Resolve a failed SCSI unconfigure operation.	“How to Resolve a Failed SCSI Unconfigure Operation” on page 81

SCSI Hot-Plugging With the `cfgadm` Command

This section describes various SCSI hot-plugging procedures that you can perform with the `cfgadm` command.

Note – The SCSI framework generally supports hot-plugging of SCSI devices. However, you should consult your hardware documentation to confirm whether hot-plugging is supported for your SCSI devices.

These procedures use specific devices as examples to illustrate how to use the `cfgadm` command to hot-plug SCSI components. The device information that you supply, and that the `cfgadm` command displays, depends on your system configuration.

All of the procedures in this section require administrative privileges that are not generally granted to user accounts. For more information, see [“How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*](#).

▼ How to Display Information About SCSI Devices

The following procedure uses SCSI controllers `c2` and `c3` and the devices that are attached to them in the examples of the type of device configuration information that you can display with the `cfgadm` command.

Note – If the SCSI device is not supported by the `cfgadm` command, the device does not display in the `cfgadm` command output.

1 Become an administrator.

2 Display information about attachment points on the system.

```
# cfgadm -l
Ap_Id                               Type           Receptacle    Occupant      Condition
c2                                   scsi-bus      connected     configured    unknown
c3                                   scsi-sas      connected     configured    unknown
```

In this example, c2 and c3 represent two SCSI controllers.

3 Display information about a system's SCSI controllers and their attached devices.

```
# cfgadm -al
Ap_Id                               Type           Receptacle    Occupant      Condition
c2                                   scsi-bus      connected     configured    unknown
c2::dsk/c2t0d0                       CD-ROM        connected     configured    unknown
c3                                   scsi-sas      connected     configured    unknown
c3::dsk/c3t0d0                       disk          connected     configured    unknown
c3::dsk/c3t1d0                       disk          connected     configured    unknown
c3::dsk/c3t2d0                       disk          connected     configured    unknown
c3::dsk/c3t3d0                       disk          connected     configured    unknown
```

Note – The `cfgadm -l` command displays information about SCSI HBAs but not SCSI devices. Use the `cfgadm -al` command to display information about SCSI devices such as disk and tapes.

▼ How to Unconfigure a SCSI Controller

The following procedure uses SCSI controller c2 in the example of unconfiguring a SCSI controller.

1 Become an administrator.**2 Unconfigure a SCSI controller.**

```
# cfgadm -c unconfigure c2
```

3 Verify that the SCSI controller is unconfigured.

```
# cfgadm -al
Ap_Id                               Type           Receptacle    Occupant      Condition
c2                                   scsi-bus      connected     unconfigured  unknown
c2::dsk/c2t0d0                       unavailable   connected     unconfigured  unknown
c3                                   scsi-sas      connected     configured    unknown
c3::dsk/c3t0d0                       disk          connected     configured    unknown
c3::dsk/c3t1d0                       disk          connected     configured    unknown
c3::dsk/c3t2d0                       disk          connected     configured    unknown
c3::dsk/c3t3d0                       disk          connected     configured    unknown
```

Notice that the `Occupant` column for c2 specifies `unconfigured`, indicating that the SCSI bus has no configured occupants.

If the unconfigure operation fails, see [“How to Resolve a Failed SCSI Unconfigure Operation”](#) on page 81.

▼ How to Configure a SCSI Controller

The following procedure uses SCSI controller `c2` in the example of configuring a SCSI controller.

1 Become an administrator.

2 Configure a SCSI controller.

```
# cfgadm -c configure c2
```

3 Verify that the SCSI controller is configured.

```
# cfgadm -al
Ap_Id                Type           Receptacle  Occupant    Condition
c2                   scsi-bus      connected   configured  unknown
c2::dsk/c2t0d0      unavailable    connected   unconfigured unknown
c3                   scsi-sas      connected   configured  unknown
c3::dsk/c3t0d0      disk          connected   configured  unknown
c3::dsk/c3t1d0      disk          connected   configured  unknown
c3::dsk/c3t2d0      disk          connected   configured  unknown
c3::dsk/c3t3d0      disk          connected   configured  unknown
```

The previous unconfigure procedure removed all devices on the SCSI bus. Now all the devices are configured back into the system.

▼ How to Configure a SCSI Device

The following procedure uses SCSI disk `c2t0d0` in the example of configuring a SCSI device.

1 Become an administrator.

2 Identify the device to be configured.

```
# cfgadm -al
Ap_Id                Type           Receptacle  Occupant    Condition
c2                   scsi-bus      connected   configured  unknown
c2::dsk/c2t0d0      unavailable    connected   unconfigured unknown
c3                   scsi-sas      connected   configured  unknown
c3::dsk/c3t0d0      disk          connected   configured  unknown
c3::dsk/c3t1d0      disk          connected   configured  unknown
c3::dsk/c3t2d0      disk          connected   configured  unknown
c3::dsk/c3t3d0      disk          connected   configured  unknown
```

3 Configure the SCSI device.

```
# cfgadm -c configure c2::dsk/c2t0d0
```

4 Verify that the SCSI device is configured.

```
# cfgadm -al
Ap_Id                Type           Receptacle  Occupant    Condition
c2                   scsi-bus      connected   configured  unknown
```

<code>c2::dsk/c2t0d0</code>	CD-ROM	connected	configured	unknown
<code>c3</code>	scsi-sas	connected	configured	unknown
<code>c3::dsk/c3t0d0</code>	disk	connected	configured	unknown
<code>c3::dsk/c3t1d0</code>	disk	connected	configured	unknown
<code>c3::dsk/c3t2d0</code>	disk	connected	configured	unknown
<code>c3::dsk/c3t3d0</code>	disk	connected	configured	unknown

▼ How to Disconnect a SCSI Controller



Caution – Disconnecting a SCSI device must be done with caution, particularly when you are dealing with controllers for disks that contain critical components of the root file system. The dynamic reconfiguration software cannot detect all cases where a system hang might result. Use this procedure with caution.

The following procedure uses SCSI controller `c2` in the example of disconnecting a SCSI device.

- 1 **Become an administrator.**
- 2 **Verify that the device is connected before you disconnect it.**

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c2                   scsi-bus            connected   configured unknown
c2::dsk/c2t0d0       CD-ROM              connected   configured unknown
c3                   scsi-sas             connected   configured unknown
c3::dsk/c3t0d0       disk                 connected   configured unknown
c3::dsk/c3t1d0       disk                 connected   configured unknown
c3::dsk/c3t2d0       disk                 connected   configured unknown
c3::dsk/c3t3d0       disk                 connected   configured unknown
```

- 3 **Disconnect the SCSI controller.**

```
# cfgadm -c disconnect c2
WARNING: Disconnecting critical partitions may cause system hang.
Continue (yes/no)? y
```



Caution – This command suspends all I/O activity on the SCSI bus until the `cfgadm -c connect` command is used. The `cfgadm` command does some basic checking to prevent critical partitions from being disconnected, but it cannot detect all cases. Inappropriate use of this command can result in a system hang and could require a system reboot.

- 4 **Verify that the SCSI bus is disconnected.**

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c2                   unavailable         disconnected  configured unknown
c2::dsk/c2t0d0       unavailable         disconnected  configured unknown
c3                   scsi-sas             connected   configured unknown
```

```

c3::dsk/c3t0d0      disk      connected  configured  unknown
c3::dsk/c3t1d0      disk      connected  configured  unknown
c3::dsk/c3t2d0      disk      connected  configured  unknown
c3::dsk/c3t3d0      disk      connected  configured  unknown

```

The controller and all the devices that are attached to it are disconnected from the system.

▼ SPARC: How to Connect a SCSI Controller

The following procedure uses SCSI controller `c2` in the example of connecting a SCSI controller.

1 Become an administrator.

2 Verify that the device is disconnected before you connect it.

```

# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c2             unavailable  disconnected  configured  unknown
c2::dsk/c2t0d0  unavailable  disconnected  configured  unknown
c3             scsi-sas     connected   configured  unknown
c3::dsk/c3t0d0  disk        connected   configured  unknown
c3::dsk/c3t1d0  disk        connected   configured  unknown
c3::dsk/c3t2d0  disk        connected   configured  unknown
c3::dsk/c3t3d0  disk        connected   configured  unknown

```

3 Connect the SCSI controller.

```
# cfgadm -c connect c2
```

4 Verify that the SCSI controller is connected.

```

# cfgadm -al
Ap_Id          Type          Receptacle  Occupant    Condition
c2             scsi-bus     connected   configured  unknown
c2::dsk/c2t0d0  CD-ROM       connected   configured  unknown
c3             scsi-sas     connected   configured  unknown
c3::dsk/c3t0d0  disk        connected   configured  unknown
c3::dsk/c3t1d0  disk        connected   configured  unknown
c3::dsk/c3t2d0  disk        connected   configured  unknown
c3::dsk/c3t3d0  disk        connected   configured  unknown

```

▼ SPARC: How to Add a SCSI Device to a SCSI Bus

SCSI controller `c3` is used in the example of how to add a SCSI device to a SCSI bus.

Note – When you add devices, you specify the `Ap_Id` of the SCSI HBA (controller) to which the device is attached, not the `Ap_Id` of the device itself.

1 Become an administrator.

2 Identify the current SCSI configuration.

```
# cfgadm -al
Ap_Id                                Type                Receptacle  Occupant  Condition
c2_                                   scsi-bus           connected   configured unknown
c2::dsk/c2t0d0                       CD-ROM             connected   configured unknown
c3_                                   scsi-sas           connected   configured unknown
c3::dsk/c3t0d0                       disk               connected   configured unknown
c3::dsk/c3t1d0                       disk               connected   configured unknown
c3::dsk/c3t2d0                       disk               connected   configured unknown
```

3 Add the SCSI device to the SCSI bus.**a. Type the following `cfgadm` command.**

For example:

```
# cfgadm -x insert_device c3
Adding device to SCSI HBA: /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2
This operation will suspend activity on SCSI bus: c3
```

b. Type `y` at the Continue (yes/no)? prompt to proceed.

```
Continue (yes/no)? y
SCSI bus quiesced successfully.
It is now safe to proceed with hotplug operation.
```

I/O activity on the SCSI bus is suspended while the hot-plug operation is in progress.

c. Connect the device and then power it on.**d. Type `y` at the Enter y if operation is complete or n to abort (yes/no)? prompt.**

```
Enter y if operation is complete or n to abort (yes/no)? y
```

4 Verify that the device has been added.

```
# cfgadm -al
Ap_Id                                Type                Receptacle  Occupant  Condition
c2_                                   scsi-bus           connected   configured unknown
c2::dsk/c2t0d0                       CD-ROM             connected   configured unknown
c3_                                   scsi-sas           connected   configured unknown
c3::dsk/c3t0d0                       disk               connected   configured unknown
c3::dsk/c3t1d0                       disk               connected   configured unknown
c3::dsk/c3t2d0                       disk               connected   configured unknown
c3::dsk/c3t3d0                       disk               connected   configured unknown
```

A new disk has been added to controller c3.

▼ SPARC: How to Replace an Identical Device on a SCSI Controller

The following procedure uses SCSI disk `c3t3d0` in the example of replacing an identical device on a SCSI controller.

Review the following conditions when attempting to replace an identical device on a SCSI controller:

- If you are replacing a mirrored or unmirrored boot device that is part of a ZFS root pool, see the following reference:

<http://www.oracle.com/technetwork/articles/servers-storage-admin/011-091-sol-dis-recovery-489183.html>

- If the device is controlled by legacy volume management software, consult your volume management documentation for the specific steps to replace a device. If you have an active Oracle support plan, see the following references:
 - If you are using Solaris Volume Manager (SVM), see MOS document 1010753.1.
 - If you are using Veritas Volume Manager (VxVM), see MOS document 1003122.1, 1011782.1, and 1002285.1.

1 Become an administrator.

2 Identify the current SCSI configuration.

```
# cfgadm -al
c2                scsi-bus    connected  configured  unknown
c2::dsk/c2t0d0    CD-ROM      connected  configured  unknown
c3                scsi-sas    connected  configured  unknown
c3::dsk/c3t0d0    disk        connected  configured  unknown
c3::dsk/c3t1d0    disk        connected  configured  unknown
c3::dsk/c3t2d0    disk        connected  configured  unknown
c3::dsk/c3t3d0    disk        connected  configured  unknown
```

3 Replace a device on the SCSI bus with another device of the same type.

a. Type the following `cfgadm` command.

For example:

```
# cfgadm -x replace_device c3::dsk/c3t3d0
Replacing SCSI device: /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@3,0
This operation will suspend activity on SCSI bus: c3
```

b. Type `y` at the Continue (yes/no)? prompt to proceed.

I/O activity on the SCSI bus is suspended while the hot-plug operation is in progress.

```
Continue (yes/no)? y
SCSI bus quiesced successfully.
It is now safe to proceed with hotplug operation.
```

c. Power off the device to be removed and remove it.

d. Add the replacement device. Then, power it on.

The replacement device should be of the same type and at the same address (target and LUN) as the device to be removed.

e. Type `y` at the Enter `y` if operation is complete or `n` to abort (yes/no)? prompt.

Enter `y` if operation is complete or `n` to abort (yes/no)? `y`

4 Verify that the device has been replaced.

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c2                   scsi-bus            connected    configured unknown
c2::dsk/c2t0d0       CD-ROM              connected    configured unknown
c3                   scsi-sas             connected    configured unknown
c3::dsk/c3t0d0       disk                 connected    configured unknown
c3::dsk/c3t1d0       disk                 connected    configured unknown
c3::dsk/c3t2d0       disk                 connected    configured unknown
c3::dsk/c3t3d0       disk                 connected    configured unknown
```

▼ SPARC: How to Remove a SCSI Device

The following procedure uses SCSI disk `c3t3d0` in the example of removing a device on a SCSI controller.

1 Become an administrator.**2 Identify the current SCSI configuration.**

```
# cfgadm -al
Ap_Id                Type                Receptacle  Occupant  Condition
c2                   scsi-bus            connected    configured unknown
c2::dsk/c2t0d0       CD-ROM              connected    configured unknown
c3                   scsi-sas             connected    configured unknown
c3::dsk/c3t0d0       disk                 connected    configured unknown
c3::dsk/c3t1d0       disk                 connected    configured unknown
c3::dsk/c3t2d0       disk                 connected    configured unknown
c3::dsk/c3t3d0       disk                 connected    configured unknown
```

3 Remove the SCSI device from the system.**a. Type the following `cfgadm` command.**

For example:

```
# cfgadm -x remove_device c3::dsk/c3t3d0
Removing SCSI device: /devices/pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@3,0
This operation will suspend activity on SCSI bus: c3
```

b. Type `y` at the Continue (yes/no)? prompt to proceed.

Continue (yes/no)? `y`
 SCSI bus quiesced successfully.
 It is now safe to proceed with hotplug operation.

I/O activity on the SCSI bus is suspended while the hot-plug operation is in progress.

c. Power off the device to be removed and remove it.

d. Type `y` at the Enter `y` if operation is complete or `n` to abort (yes/no)? prompt.Enter `y` if operation is complete or `n` to abort (yes/no)? `y`

Note – This step must be performed if you are removing a SCSI RAID device from a SCSI RAID array.

4 Verify that the device has been removed from the system.

```
# cfgadm -al
Ap_Id                Type      Receptacle  Occupant  Condition
c2                   scsi-bus connected   configured unknown
c2:::dsk/c2t0d0      CD-ROM    connected   configured unknown
c3                   scsi-sas  connected   configured unknown
c3:::dsk/c3t0d0      disk      connected   configured unknown
c3:::dsk/c3t1d0      disk      connected   configured unknown
c3:::dsk/c3t2d0      disk      connected   configured unknown
```

Troubleshooting SCSI Configuration Problems

This section provides error messages and possible solutions for troubleshooting SCSI configuration problems. For more information on troubleshooting SCSI configuration problems, see [`cfgadm\(1M\)`](#).

Error Message

```
cfgadm: Component system is busy, try again: failed to offline:
  device-path
    Resource          Information
-----
/dev/dsk/c1t0d0s0    mounted filesystem "file-system"
```

Cause

You attempted to remove or replace a device with a mounted file system.

Solution

Unmount the file system that is listed in the error message and retry the `cfgadm` operation.

Error Message

```
cfgadm: Component system is busy, try again: failed to offline:
  device-path
    Resource          Information
-----
/dev/dsk/device-name  swap area
```

Cause

If you use the `cfgadm` command to remove a system resource, such as a swap device or a dedicated dump device, a similar error message is displayed if the system resource is still active.

Solution

Unconfigure the swap areas on the device that is specified and retry the `cfgadm` operation.

Error Message

```

cfgadm: Component system is busy, try again: failed to offline:
  device-path
    Resource           Information
  -----
/dev/dsk/device-name  dump device (swap)

```

Cause

You attempted to remove or replace a dump device that is configured on a swap area.

Solution

Unconfigure the dump device that is configured on the swap area and retry the `cfgadm` operation.

Error Message

```

cfgadm: Component system is busy, try again: failed to offline:
  device-path
    Resource           Information
  -----
/dev/dsk/device-name  dump device (dedicated)

```

Cause

You attempted to remove or replace a dedicated dump device.

Solution

Unconfigure the dedicate dump device and retry the `cfgadm` operation.

▼ How to Resolve a Failed SCSI Unconfigure Operation

Use this procedure if one or more target devices are busy and the SCSI unconfigure operation fails. Otherwise, future dynamic reconfiguration operations on this controller and target devices will fail with a `dr in progress` message.

- 1 **Become an administrator.**
- 2 **Reconfigure the controller.**

```
# cfgadm -c configure device-name
```

PCI or PCIe Hot-Plugging With the `cfgadm` Command (Task Map)

The following task map describes the tasks for managing PCI or PCIe devices on your system.

Task	Description	For Instructions
Display PCI slot configuration information.	Display the status of PCI hot-pluggable devices and slots on the system.	“How to Display PCI Slot Configuration Information” on page 83
Remove a PCI adapter card.	Unconfigure the card, disconnect power from the slot, and remove the card from the system.	“How to Remove a PCI Adapter Card” on page 84
Add a PCI adapter card.	Insert the adapter card into a hot-pluggable slot, connect power to the slot, and configure the card.	“How to Add a PCI Adapter Card” on page 86
Troubleshoot PCI configuration problems.	Identify error message and possible solutions to resolve PCI configuration problems.	“Troubleshooting PCI Configuration Problems” on page 87

PCI or PCIe Hot-Plugging With the `cfgadm` Command

This section provides step-by-step instructions for hot-plugging PCI or PCIe adapter cards on SPARC and x86 based systems.

In addition to the `cfgadm` command, the `prtconf` command is helpful during hot-plug operations. The `prtconf` command displays additional configuration information that pertains to the hardware.

After adding hardware, you can use the `prtconf` command to verify that the hardware is configured correctly. For example, after a configure operation, use the `prtconf -D` command to verify that the driver is attached to the newly installed hardware device. If the device driver has not been added to the system prior to hardware configuration, it may be manually added by using the `add_drv` command.

For more information, see [`prtconf\(1M\)`](#) and [`add_drv\(1M\)`](#).

All of the procedures in this section require administrative privileges that are not generally granted to user accounts. For more information, see [“How to Use Your Assigned Administrative Rights”](#) in *Oracle Solaris 11.1 Administration: Security Services*.

In the examples, only PCI attachment points are listed, for brevity. The attachment points that are displayed on your system depend on your system configuration.

PCIe LED Indicator Behavior

You might observe the LED indicators on the system to get a visual indication about the status of the slot's hot-plug operation. The LED behavior, in case of PCI Express, matches that defined in the PCI Express specification or the behavior might otherwise be platform dependent.

Please refer to your platform guide for specific details. In case of PCI Express, when the Attention Button is pressed, the power indicator blinks, which indicates the beginning of a state transition. The blinking ends when the state transition has ended.

▼ How to Display PCI Slot Configuration Information

This procedure has been updated to include PCIe configuration information.

The `cfgadm` command displays the status of PCI hot-pluggable devices and slots on a system. For more information, see [`cfgadm\(1M\)`](#).

- 1 **Become an administrator.**
- 2 **Display PCI configuration information.**
 - **Display PCI slot configuration information.**

For example:

```
# cfgadm
Ap_Id                Type                Receptacle  Occupant  Condition
pci1:hpc0_slot0     unknown            empty       unconfigured unknown
pci1:hpc0_slot1     unknown            empty       unconfigured unknown
pci1:hpc0_slot2     unknown            empty       unconfigured unknown
pci1:hpc0_slot3     ethernet/hp        connected   configured ok
pci1:hpc0_slot4     unknown            empty       unconfigured unknown
```

- **Display specific PCI device information.**

For example:

```
# cfgadm -s "cols=ap_id:type:info" pci
Ap_Id                Type                Information
pci1:hpc0_slot0     unknown            Slot 7
pci1:hpc0_slot1     unknown            Slot 8
pci1:hpc0_slot2     unknown            Slot 9
pci1:hpc0_slot3     ethernet/hp        Slot 10
pci1:hpc0_slot4     unknown            Slot 11
```

The logical `Ap_Id`, `pci1:hpc0_slot0`, is the logical `Ap_Id` for hot-pluggable slot, Slot 7. The component `hpc0` indicates the hot-pluggable adapter card for this slot, and `pci1` indicates the PCI bus instance. The `Type` field indicates the type of PCI adapter card that is present in the slot.

- **Display PCIe slot configuration information.**

For example:

```
# cfgadm pci
Ap_Id                Type                Receptacle  Occupant  Condition
pcie1                unknown            empty       unconfigured unknown
pcie2                unknown            empty       unconfigured unknown
pcie3                unknown            empty       unconfigured unknown
```

pcie4	etherne/hp	connected	configured	ok
pcie5	pci-pci/hp	connected	configured	ok
pcie6	unknown	disconnected	unconfigured	unknown

- **Display specific PCIe device information.**

For example:

```
# cfgadm -s "cols=ap_id:busy:o_state" pci
Ap_Id      Busy      Occupant
pcie1      n         unconfigured
pcie2      n         unconfigured
pcie3      n         unconfigured
pcie4      n         configured
pcie5      n         configured
pcie6      n         configured
```

Note – The logical Ap_Id in most cases should match the slot label that is silk-screened on the system chassis. Refer to your platform guide for the cfgadm output of the hot-pluggable slots. The Busy field can be displayed to ensure that the Ap_Id is not transitioning to another state before a hot-plug operation is attempted.

▼ How to Remove a PCI Adapter Card

The following procedure has been updated for removing a PCIe adapter card. However, the procedure to remove an adapter card is the same whether you are using PCI or PCIe.

- 1 **Become an administrator.**
- 2 **Determine which slot the PCI adapter card is in.**

For example:

```
# cfgadm pci
Ap_Id      Type      Receptacle  Occupant    Condition
pcie1      unknown   empty        unconfigured unknown
pcie2      unknown   empty        unconfigured unknown
pcie3      unknown   empty        unconfigured unknown
pcie4      etherne/hp  connected    configured  ok
pcie5      pci-pci/hp  connected    configured  ok
pcie6      unknown   disconnected  unconfigured unknown
```

- 3 **Stop the application that has the device open.**

For example, if the device is an Ethernet card, use the ipadm command to bring down the interface and unplumb the interface. For example:

```
# ipadm delete-ip bge1
```

- 4 Unconfigure the device manually by using the `cfgadm` command as shown below. Or, if you have a PCIe adapter card, use the auto-configuration method, such as pressing the slot's Attention Button as defined by your platform guide.**

```
# cfgadm -c unconfigure pcie4
```

- 5 Confirm that the device has been unconfigured.**

For example:

```
# cfgadm pci
Ap_Id                Type      Receptacle  Occupant    Condition
pcie1                unknown  empty       unconfigured unknown
pcie2                unknown  empty       unconfigured unknown
pcie3                unknown  empty       unconfigured unknown
pcie4                unknown  connected   unconfigured unknown
pcie5                pci-pci/hp connected  configured  ok
pcie6                unknown  disconnected unconfigured unknown
```

Note – The `Type` and `Condition` also become unknown when the device is unconfigured.

- 6 Disconnect the power to the slot manually. If the auto-configuration method is used, this step is not necessary.**

```
# cfgadm -c disconnect pcie4
```

Refer to your platform guide for more information.

- 7 Confirm that the device has been disconnected.**

For example:

```
# cfgadm pci
Ap_Id                Type      Receptacle  Occupant    Condition
pcie1                unknown  empty       unconfigured unknown
pcie2                unknown  empty       unconfigured unknown
pcie3                unknown  empty       unconfigured unknown
pcie4                unknown  disconnected unconfigured unknown
pcie5                pci-pci/hp connected  configured  ok
pcie6                unknown  disconnected unconfigured unknown
```

- 8 Follow appropriate instructions in your platform guide to remove the PCI adapter card. After the card is removed, the `Receptacle` state is empty.**

For example:

```
# cfgadm pci
Ap_Id                Type      Receptacle  Occupant    Condition
pcie1                unknown  empty       unconfigured unknown
pcie2                unknown  empty       unconfigured unknown
pcie3                unknown  empty       unconfigured unknown
pcie4                unknown  empty       unconfigured unknown
pcie5                pci-pci/hp connected  configured  ok
pcie6                unknown  disconnected unconfigured unknown
```

Note – The auto-configuration method can be enabled or disabled at boot, depending on the platform implementation. Set the auto-configuration method as appropriate for your environment.

▼ How to Add a PCI Adapter Card

The following procedure has been updated for adding a PCIe adapter card. However, the procedure to add an adapter card is the same whether you are using PCI or PCIe.

1 Become an administrator.

2 Identify the hot-pluggable slot and open latches.

For example, `pcie3`.

```
# cfgadm pci
Ap_Id          Type          Receptacle  Occupant    Condition
pcie1          unknown      empty        unconfigured unknown
pcie2          unknown      empty        unconfigured unknown
pcie3          unknown      empty        unconfigured unknown
pcie4          unknown      empty        unconfigured unknown
pcie5          pci-pci/hp   connected    configured  ok
pcie6          unknown      disconnected  unconfigured unknown
```

3 Follow the appropriate instructions in your platform guide to insert a PCI adapter card into the slot.

4 Determine which slot the PCI adapter card is in after it is inserted.

For example:

```
# cfgadm pci
Ap_Id          Type          Receptacle  Occupant    Condition
pcie1          unknown      empty        unconfigured unknown
pcie2          unknown      empty        unconfigured unknown
pcie3          unknown      disconnected  unconfigured unknown
pcie4          unknown      empty        unconfigured unknown
pcie5          pci-pci/hp   connected    configured  ok
pcie6          unknown      disconnected  unconfigured unknown
```

5 Connect the power to the slot manually using the `cfgadm` command. Or, if you have a PCIe adapter card, use the auto-configuration method, such as pressing the slot's Attention Button as defined by your platform guide.

For example:

```
# cfgadm -c connect pcie3
```

6 Confirm that the attachment point is connected.

For example:

```
# cfgadm pci
Ap_Id          Type          Receptacle  Occupant    Condition
pcie1          unknown      empty        unconfigured unknown
pcie2          unknown      empty        unconfigured unknown
pcie3          unknown      connected    unconfigured unknown
pcie4          unknown      empty        unconfigured unknown
pcie5          pci-pci/hp   connected    configured   ok
pcie6          unknown      disconnected  unconfigured unknown
```

7 Configure the PCI adapter card manually by using the `cfgadm` command as shown below. If using the auto-configuration method, this step should not be necessary. Refer to your platform guide for more information.

For example:

```
# cfgadm -c configure pcie3
```

8 Verify the configuration of the PCI adapter card in the slot.

For example:

```
# cfgadm pci
Ap_Id          Type          Receptacle  Occupant    Condition
pcie1          unknown      empty        unconfigured unknown
pcie2          unknown      empty        unconfigured unknown
pcie3          etherne/hp   connected    configured   unknown
pcie5          pci-pci/hp   connected    configured   ok
pcie6          unknown      disconnected  unconfigured unknown
```

9 Configure any supporting software if this device is a new device.

For example, if this device is an Ethernet card, use the `ipadm` command to set up the interface.

For example:

```
# ipadm create-addr bge1
```

Note – The auto-configuration method can be enabled or disabled at boot, depending on the platform implementation. Set the auto-configuration method as appropriate for your environment.

Troubleshooting PCI Configuration Problems

Error Message

```
cfgadm: Configuration operation invalid: invalid transition
```

Cause

An invalid transition was attempted.

Solution

Check whether the `cfgadm -c` command was issued appropriately. Use the `cfgadm` command to check the current receptacle and occupant state and to make sure that the `Ap_Id` is correct.

Error Message

```
cfgadm: Attachment point not found
```

Cause

The specified attachment point was not found.

Solution

Check whether the attachment point is correct. Use the `cfgadm` command to display a list of available attachment points. Also check the physical path to see if the attachment point is still there.

SATA Hot-Plugging With the `cfgadm` Command

SATA controller and port multiplier device ports are represented by attachment points in the device tree. SATA devices that are connected and configured on the system are shown as attachment point name extension. The terms attachment point and SATA port can be used interchangeably.

The `cfgadm` syntax that is used with SATA devices is slightly different than `cfgadm` syntax for SCSI or PCI devices.

You can display SATA device information as follows:

```
% cfgadm -al
Ap_Id                               Type      Receptacle  Occupant    Condition
sata0/0::dsk/c7t0d0                 disk      connected   configured  ok
sata0/1::dsk/c7t1d0                 disk      connected   configured  ok
sata0/2::dsk/c7t2d0                 disk      connected   configured  ok
sata0/3::dsk/c7t3d0                 disk      connected   configured  ok
.
.
.
```

▼ How to Unconfigure a SATA Device

In general, SATA devices must be unconfigured before they can be removed and replaced. If you attempt to unconfigure a device that is part of an active ZFS storage pool, you will see an error message similar to the following:

```
# cfgadm -c unconfigure sata5/7
Unconfigure the device at: /devices/pci@2,0/pci1022...
This operation will suspend activity on the SATA device
```

Continue (yes/no)? y

cfgadm: Hardware specific failure: Failed to unconfig device at ap_id: /devices/pci@2,0/pci10...

1 Become an administrator.

2 Identify the device to be unconfigured.

```
# cfgadm -al | grep c7t0d0
sata0/0::dsk/c7t0d0          disk          connected   configured  ok
```

3 Unconfigure the device.

```
# cfgadm -c unconfigure sata0/0
```

If you attempt to unconfigure the device by specifying the individual device, you will see a message similar to the following:

```
# cfgadm -c unconfigure sata0/0::dsk/c7t0d0
do_control_ioctl: open failed: errno:2
cfgadm: Library error: Cannot open ap_id: /devices/pci@0,0/pci10...
No such file or directory
```

4 Confirm that the device is unconfigured.

```
# cfgadm | grep sata0/0
sata0/0          disk          connected   unconfigured ok
```

▼ How to Configure a SATA Device

After a disk is physically removed or replaced, it can be configured.

1 Become an administrator.

2 Configure the device.

```
# cfgadm -c configure sata0/0
```

3 Confirm that the device is configured.

```
# cfgadm | grep sata0/0
sata0/0::dsk/c7t0d0          disk          connected   configured  ok
```

Reconfiguration Coordination Manager (RCM) Script Overview

The Reconfiguration Coordination Manager (RCM) is the framework that manages the dynamic removal of system components. By using RCM, you can register and release system resources in an orderly manner.

You can use the new RCM script feature to write your own scripts to shut down your applications, or to cleanly release the devices from your applications during dynamic reconfiguration. The RCM framework launches a script automatically in response to a reconfiguration request, if the request impacts the resources that are registered by the script.

You can also release resources from applications manually before you dynamically remove the resource. Or, you can use the `cfgadm` command with the `-f` option to force a reconfiguration operation. However, this option might leave your applications in an unknown state. Also, the manual release of resources from applications commonly causes errors.

The RCM script feature simplifies and better controls the dynamic reconfiguration process. By creating an RCM script, you can do the following:

- Automatically release a device when you dynamically remove a device. This process also closes the device if the device is opened by an application.
- Run site-specific tasks when you dynamically remove a device from the system.

What Is an RCM Script?

- An executable shell script (Perl, sh, csh, or ksh) or binary program that the RCM daemon runs. Perl is the recommended language.
- A script that runs in its own address space by using the user ID of the script file owner.
- A script that is run by the RCM daemon when you use the `cfgadm` command to dynamically reconfigure a system resource.

What Can an RCM Script Do?

You can use an RCM script to release a device from an application when you dynamically remove a device. If the device is currently open, the RCM script also closes the device.

For example, an RCM script for a tape backup application can inform the tape backup application to close the tape drive or shut down the tape backup application.

How Does the RCM Script Process Work?

You can invoke an RCM script as follows:

```
$ script-name command [args ...]
```

An RCM script performs the following basic steps:

1. Takes the RCM command from command-line arguments.

2. Executes the command.
3. Writes the results to `stdout` as name-value pairs.
4. Exits with the appropriate exit status.

The RCM daemon runs one instance of a script at a time. For example, if a script is running, the RCM daemon does not run the same script until the first script exits.

RCM Script Commands

You must include the following RCM commands in an RCM script:

- `scriptinfo` – Gathers script information
- `register` – Registers interest in resources
- `resourceinfo` – Gathers resource information

You might include some or all of the following RCM commands:

- `queryremove` – Queries whether the resource can be released
- `preremove` – Releases the resource
- `postremove` – Provides post-resource removal notification
- `undoremove` – Undoes the actions done in `preremove`

For a complete description of these RCM commands, see [rcmscript\(4\)](#).

RCM Script Processing Environment

When you dynamically remove a device, the RCM daemon runs the following:

- The script's `register` command to gather the list of resources (device names) that are identified in the script.
- The script's `queryremove` and `preremove` commands prior to removing the resource if the script's registered resources are affected by the dynamic remove operation.
- The script's `postremove` command if the remove operation succeeds. However, if the remove operation fails, the RCM daemon runs the script's `undoremove` command.

RCM Script Tasks

The following sections describe the RCM script tasks for application developers and system administrators.

Application Developer RCM Script (Task Map)

The following task map describes the tasks for an application developer who is creating an RCM script.

Task	Description	For Instructions
1. Identify the resources your application uses.	Identify the resources (device names) your application uses that you could potentially dynamically remove.	cfgadm(1M)
2. Identify the commands to release the resource.	Identify the commands for notifying the application to cleanly release the resource from the application.	Application documentation
3. Identify the commands for post-removal of the resource.	Include the commands for notifying the application of the resource removal.	rcmscript(4)
4. Identify the commands if the resource removal fails.	Include the commands for notifying the application of the available resource.	rcmscript(4)
5. Write the RCM script.	Write the RCM script based on the information identified in tasks 1-4.	“Tape Backup RCM Script Example” on page 95
6. Install the RCM script.	Add the script to the appropriate script directory.	“How to Install an RCM Script” on page 93
7. Test the RCM script	Test the script by running the script commands manually and by initiating a dynamic reconfiguration operation.	“How to Test an RCM Script” on page 94

System Administrator RCM Script (Task Map)

The following task map describes the tasks for a system administrator who is creating an RCM script to do site customization.

Task	Description	For Instructions
1. Identify the resources to be dynamically removed.	Identify the resources (device names) to be potentially removed by using the <code>cfgadm -l</code> command.	cfgadm(1M)
2. Identify the applications to be stopped.	Identify the commands for stopping the applications cleanly.	Application documentation
3. Identify the commands for pre-removal and post-removal of the resource.	Identify the actions to be taken before and after the resource is removed.	rcmscript(4)

Task	Description	For Instructions
4. Write the RCM script.	Write the RCM script based on the information identified in tasks 1-3.	“Tape Backup RCM Script Example” on page 95
5. Install the RCM script.	Add the script to the appropriate script directory.	“How to Install an RCM Script” on page 93
6. Test the RCM script.	Test the script by running the script commands manually and by initiating a dynamic reconfiguration operation.	“How to Test an RCM Script” on page 94

Naming an RCM Script

A script must be named as *vendor,service* where the following applies:

vendor Is the stock symbol of the vendor that provides the script, or any distinct name that identifies the vendor.

service Is the name of the service that the script represents.

Installing or Removing an RCM Script

You must be an administrator to install or remove an RCM script. Use this table to determine where you should install your RCM script.

TABLE 4-1 RCM Script Directories

Directory Location	Script Type
/etc/rcm/scripts	Scripts for specific systems
/usr/platform/'uname -i'/lib/rcm/scripts	Scripts for a specific hardware implementation
/usr/platform/'uname -m'/lib/rcm/scripts	Scripts for a specific hardware class
/usr/lib/rcm/scripts	Scripts for any hardware

▼ How to Install an RCM Script

- 1 **Become an administrator.**
- 2 **Copy the script to the appropriate directory.**
See [Table 4-1](#).

For example:

```
# cp ABC,sample.pl /usr/lib/rcm/scripts
```

- 3 **Change the user ID and the group ID of the script to the desired values.**

```
# chown user:group /usr/lib/rcm/scripts/ABC,sample.pl
```

- 4 **Send SIGHUP to the RCM daemon.**

```
# pkill -HUP -x -u root rcm_daemon
```

▼ How to Remove an RCM Script

- 1 **Become an administrator.**

- 2 **Remove the script from the RCM script directory.**

For example:

```
# rm /usr/lib/rcm/scripts/ABC,sample.pl
```

- 3 **Send SIGHUP to the RCM daemon.**

```
# pkill -HUP -x -u root rcm_daemon
```

▼ How to Test an RCM Script

- 1 **Set environment variables, such as RCM_ENV_FORCE, in the command-line shell before running your script.**

For example, in the Korn shell, use the following:

```
$ export RCM_ENV_FORCE=TRUE
```

- 2 **Test the script by running the script commands manually from the command line.**

For example:

```
$ script-name scriptinfo
$ script-name register
$ script-name preremove resource-name
$ script-name postremove resource-name
```

- 3 **Make sure that each RCM script command in your script prints appropriate output to stdout.**

- 4 **Install the script in the appropriate script directory.**

For more information, see [“How to Install an RCM Script”](#) on page 93.

5 Test the script by initiating a dynamic remove operation.

For example, assume your script registers the device, `/dev/dsk/c1t0d0s0`. Try these commands.

```
$ cfgadm -c unconfigure c1::dsk/c1t0d0
$ cfgadm -f -c unconfigure c1::dsk/c1t0d0
$ cfgadm -c configure c1::dsk/c1t0d0
```



Caution – Make sure that you are familiar with these commands because they can alter the state of the system and cause system failures.

Tape Backup RCM Script Example

This example illustrates how to use an RCM script for tape backups.

What the Tape Backup RCM Script Does

The tape backup RCM script performs the following steps:

1. Sets up a dispatch table of RCM commands.
2. Calls the dispatch routine that corresponds to the specified RCM command and exits with status 2 for unimplemented RCM commands.
3. Sets up the `scriptinfo` section.

```
rcm_script_func_info=Tape backup appl script for DR
```

4. Registers all tape drives in the system by printing all tape drive device names to `stdout`.

```
rcm_resource_name=/dev/rmt/$f
```

If an error occurs, the script prints the error information to `stdout`.

```
rcm_failure_reason=$errmsg
```

5. Sets up the resource information for the tape device.

```
rcm_resource_usage_info=Backup Tape Unit Number $unit
```

6. Sets up the `preremove` information by checking if the backup application is using the device. If the backup application is not using the device, the dynamic reconfiguration operation continues. If the backup application is using the device, the script checks `RCM_ENV_FORCE`. If `RCM_ENV_FORCE` is set to `FALSE`, the script denies the dynamic reconfiguration operation and prints the following message:

```
rcm_failure_reason=tape backup in progress pid=...
```

If `RCM_ENV_FORCE` is set to `TRUE`, the backup application is stopped, and the reconfiguration operation proceeds.

Outcomes of the Tape Backup Reconfiguration Scenarios

Here are the various outcomes if you use the `cfgadm` command to remove a tape device without the RCM script:

- If you use the `cfgadm` command and the backup application is not using the tape device, the operation succeeds.
- If you use the `cfgadm` command and the backup application is using the tape device, the operation fails.

Here are the various outcomes if you use the `cfgadm` command to remove a tape device with the RCM script.

- If you use the `cfgadm` command and the backup application is not using the tape device, the operation succeeds.
- If you use the `cfgadm` command without the `-f` option and the backup application is using the tape device, the operation fails with an error message similar to the following:


```
tape backup in progress pid=...
```
- If you use the `cfgadm -f` command and the backup application is using the tape device, the script stops the backup application and the `cfgadm` operation succeeds.

Example—Tape Backup RCM Script

```
#!/usr/bin/perl -w
#
# A sample site customization RCM script.
#
# When RCM_ENV_FORCE is FALSE this script indicates to RCM that it cannot
# release the tape drive when the tape drive is being used for backup.
#
# When RCM_ENV_FORCE is TRUE this script allows DR removing a tape drive
# when the tape drive is being used for backup by killing the tape
# backup application.
#

use strict;

my ($cmd, %dispatch);
$cmd = shift(@ARGV);
# dispatch table for RCM commands
%dispatch = (
    "scriptinfo"    =>    \do_scriptinfo,
    "register"      =>    \do_register,
    "resourceinfo" =>    \do_resourceinfo,
    "queryremove"  =>    \do_preremove,
    "preremove"    =>    \do_preremove
);

if (defined($dispatch{$cmd})) {
    &{$dispatch{$cmd}};
} else {
```

```

        exit (2);
    }

    sub do_scriptinfo
    {
        print "rcm_script_version=1\n";
        print "rcm_script_func_info=Tape backup appl script for DR\n";
        exit (0);
    }

    sub do_register
    {
        my ($dir, $f, $errmsg);

        $dir = opendir(RMT, "/dev/rmt");
        if (!$dir) {
            $errmsg = "Unable to open /dev/rmt directory: $!";
            print "rcm_failure_reason=$errmsg\n";
            exit (1);
        }

        while ($f = readdir(RMT)) {
            # ignore hidden files and multiple names for the same device
            if (($f !~ /\^\.\/) && ($f =~ /^[0-9]+$/)) {
                print "rcm_resource_name=/dev/rmt/$f\n";
            }
        }

        closedir(RMT);
        exit (0);
    }

    sub do_resourceinfo
    {
        my ($rsrc, $unit);

        $rsrc = shift(@ARGV);
        if ($rsrc =~ /\^\/dev\/rmt\/([0-9]+)$/) {
            $unit = $1;
            print "rcm_resource_usage_info=Backup Tape Unit Number $unit\n";
            exit (0);
        } else {
            print "rcm_failure_reason=Unknown tape device!\n";
            exit (1);
        }
    }

    sub do_preremove
    {
        my ($rsrc);

        $rsrc = shift(@ARGV);

        # check if backup application is using this resource
        #if (the backup application is not running on $rsrc) {
            # allow the DR to continue
        #    exit (0);
        #}
        #
    }

```

```
# If RCM_ENV_FORCE is FALSE deny the operation.
# If RCM_ENV_FORCE is TRUE kill the backup application in order
# to allow the DR operation to proceed
#
if ($ENV{RCM_ENV_FORCE} eq 'TRUE') {
    if ($cmd eq 'preremove') {
        # kill the tape backup application
    }
    exit (0);
} else {
    #
    # indicate that the tape drive can not be released
    # since the device is being used for backup by the
    # tape backup application
    #
    print "rcm_failure_reason=tape backup in progress pid=...\n"
;
    exit (3);
}
}
```

Managing USB Devices (Tasks)

This chapter provides overview information and step-by-step instructions for using USB devices in the Oracle Solaris OS. This chapter specifically covers how to use USB devices with HAL services.

The following is a list of the information in this chapter:

- “What's New in USB Devices?” on page 99
- “About USB Support in Oracle Solaris” on page 100
- “Overview of USB Devices” on page 105
- “Managing USB Mass Storage Devices” on page 109
- “Using USB Audio Devices” on page 126
- “Hot-Plugging USB Devices With the `cfgadm` Command” on page 129

For general information about dynamic reconfiguration and hot-plugging, see [Chapter 4, “Dynamically Configuring Devices \(Tasks\)”](#).

For information on configuring USB printers, see [Chapter 1, “Setting Up and Administering Printers by Using CUPS \(Overview\)”](#) in *Configuring and Managing Printing in Oracle Solaris 11.1*.

What's New in USB Devices?

The following features are new or changed in this release.

USB 3.0 Support

USB 3.0 support is provided by the introduction of a new USB host controller driver, `xhci`, the enhancement of the hub driver, and the Oracle Solaris USB Architecture (USBA) framework. USB 3.0 hub and mass-storage devices can now work in USB 3.0 mode, when they are inserted

into USB 3.0 ports. In addition, the legacy USB devices continue to work when they are connected to USB 3.0 ports, with the exception of USB audio devices. See [“Supported USB Features”](#) on page 100.

About USB Support in Oracle Solaris

The following information is described in this section:

- [“Supported USB Features”](#) on page 100
- [“USB Device Features and Compatibility Issues”](#) on page 101
- [“Bus-Powered Devices”](#) on page 102
- [“USB Keyboards and Mouse Devices”](#) on page 102
- [“USB Host Controller and Hubs”](#) on page 103
- [“USB Hub Devices”](#) on page 104
- [“SPARC: USB Power Management”](#) on page 104
- [“Guidelines for Connecting USB Cables”](#) on page 105

Supported USB Features

This Oracle Solaris release includes expanded supported for USB 3.0 devices (“SuperSpeed” devices) that use the USB 3.0 xhci host controller driver. The xhci host controller driver supports all USB devices, with the exception of *audio devices*. The xhci host controller driver supports interrupt, control, and bulk transfers for USB hosts, including super speed and non-super speed bus interfaces. Compatibility with USB 2.0, 1.1, and 1.0 devices and drivers is supported so that you can use the same cables, connectors, and software interfaces.

Support for the following USB devices is provided on both SPARC and x86 based systems, except where indicated:

- USB 3.0 devices work in USB 3.0 mode when they are inserted into USB 3.0 ports. Legacy USB devices (2.0, 1.1, and 1.0) continue to work when they are connected to USB 3.0 ports, with the exception of USB audio devices.
- Audio device support is provided for USB 2.0, 1.1, and 1.0 audio devices *only*.
Devices that are not supported by a USB driver might include libusb applications, such as `gphoto2`, `gtkam`, and `pilotlink`. For more information, refer to the user-space USB library documentation in `/usr/share/doc/libusb/libusb.txt`.
- Generic USB driver support – See [ugen\(7D\)](#).
- Human Interface Device (HID) support (keyboard and mouse devices) – See [hid\(7D\)](#)
- Hub support – See [hubd\(7D\)](#).
- Printer support
- USB CDC-ECM (USB Communication Device Class- Ethernet Control Model) device support

- Serial device support, including the following:
 - Edgeport – See [usbser_edge\(7D\)](#).
 - Prolific – See [usbspr1\(7D\)](#).
 - Keyspan – See [usbsksp\(7D\)](#).
- Storage device support – See [scsa2usb\(7D\)](#).
- User-space USB device management library support. See [libusb\(3LIB\)](#).
- Increased USB 3.0 bus speed from 480 Mbits/sec to 5 Gbits/sec.

A USB 3.0 or USB 2.0 port might be one of the following possibilities:

- A port on a USB 3.0 or USB 2.0 PCI card
- A port on a USB 3.0 or USB 2.0 hub that is connected to USB 3.0 or USB 2.0 port
- A port on a SPARC or x86 computer motherboard

Note – A USB 2.0 PCI card might be needed for older SPARC platforms.

- USB 2.0 and USB 1.1 devices work as they have in the past, even if you have USB 3.0, USB 2.0, and USB 1.0 devices on the same system. A USB 2.0, and a USB 1.0 host controller supports all four transfer types: control, bulk, interrupt, and isochronous. A USB 3.0 host controller supports three transfer types: control, bulk, and interrupt. The isochronous transfer type is not supported under a USB 3.0 host controller.
- While USB 3.0 and USB 2.0 devices can operate on a USB 1.1 or 1.0 port, their performance is significantly better when they are connected to a USB 2.0 or a USB 3.0 port.
- An XHCI host controller does not have a companion controllers, meaning a single controller supports all speeds of the USB devices. A USB 2.0 host controller has one high-speed EHCI and one or more OHCI or UHCI embedded controllers. Devices that are connected to a USB 2.0 port are dynamically assigned to either a EHCI or OHCI controller, depending on whether they support USB 2.0.

For more information about USB 3.0 device support, see [xhci\(7D\)](#).

For more information on USB 2.0 device support, see [ehci\(7D\)](#) and [usba\(7D\)](#).

For more information about USB 1.1 device support, see [ohci\(7D\)](#).

For a description of USB devices and terminology, see “Overview of USB Devices” on page 105.

USB Device Features and Compatibility Issues

USB 3.0 devices are defined as super speed (“SuperSpeed”) devices that follow the USB 3.0 specification. USB 2.0 devices are defined as high-speed (“Hi-Speed”) devices that follow the USB 2.0 specification. You can refer to the USB 3.0 and the USB 2.0 specifications at <http://www.usb.org/home>.

To identify the speed of your USB device, check the `/var/adm/messages` file for messages similar to the following:

```
Dec 13 17:05:57 mysystem usba: [ID 912658 kern.info] USB 2.0 device
(usb50d,249) operating at hi speed (USB 2.x) on USB 2.0 external hub:
storage@4, scsa2usb0 at bus address 4
```

Here are some of the USB devices that are supported in this Oracle Solaris release:

- Mass storage devices, such as CD-RWs, hard disks, DVDs, digital cameras, diskettes, tape drives, memory sticks, and multi-format card readers
- Keyboards and mouse devices
- Audio devices, such as speakers and microphones

Note – Audio device support is not available in USB 3.0.

Additional storage devices might work by modifying the `scsa2usb.conf` file. For more information, see [scsa2usb\(7D\)](#).

Bus-Powered Devices

Bus-powered hubs use power from the USB bus to which they are connected, to power devices connected to them. Special care must be taken to not overload these hubs, because the power these hubs offer to their downstream devices is limited.

Power budgeting is implemented for USB devices. This feature has the following limitations:

- Cascading two bus-powered hubs is not recommended.
- Each bus-powered hub is allowed a maximum of 100 mA only for each port.
- Only self-powered or low bus-powered devices are allowed to connect to a bus-powered hub. High bus-powered devices are denied the connection. Some hubs or devices can report a false power source, such that the connection might be unpredictable.

USB Keyboards and Mouse Devices

Keep the following issues in mind when using USB keyboards and mouse devices:

- *Do not* move the keyboard and mouse *during* a reboot or at the ok prompt on a SPARC system. You can move the keyboard and mouse to another hub at any time *after* a system reboot. After you plug in a keyboard and mouse, they are fully functional again.
- The keys just to the left of the keypad might not function on some third-party USB keyboards.

- **SPARC** – Keep the following issues in mind when using USB keyboards and mouse devices on SPARC systems:
 - The power key on a USB keyboard behaves differently than the power key on the type 5 keyboard. On a USB keyboard, you can suspend or shut down the system by using the SUSPEND/SHUTDOWN key. However, you cannot use that key to power up the system.
 - Before the boot process finishes, the OpenBoot PROM (OBP) limits keyboard and mouse devices to the motherboard root hub ports only.
 - USB keyboard and mouse devices cannot be used simultaneously with Type 3, 4, or 5 keyboards on legacy SPARC systems.
- For information about multiple keyboard and mouse device support, see [virtualkm\(7D\)](#).

USB Wheel Mouse Support

The following wheel mouse features are supported:

- Support for more than 3 buttons is available on USB or PS/2 mouse devices.
- Wheel mouse scrolling is available on a USB or PS/2 mouse device. This support means that rolling the wheel on a USB or a PS/2 mouse results in a scroll in the application or window under mouse focus. StarOffice, Firefox, and GNOME applications support wheel mouse scrolling. However, other applications might not support this functionality.

USB Host Controller and Hubs

A USB hub is responsible for the following tasks:

- Monitoring the insertion or removal of a device on its ports
- Power managing individual devices on its ports
- Controlling power to its ports

The USB host controller has an embedded hub, called the *root hub*. The ports that are visible at the system's back panel are the ports of the root hub.

The USB host controller is responsible for the following tasks:

- Directing the USB bus. Individual devices cannot arbitrate for the bus.
- Polling the devices by using a polling interval that is determined by the device. The device is assumed to have sufficient buffering to account for the time between the polls.
- Sending data between the USB host controller and its attached devices. Peer-to-peer communication is not supported.

USB Hub Devices

Keeping the following key points in mind when using USB hub devices:

- Do not cascade hubs beyond four levels on either SPARC based systems or x86 based systems. On SPARC systems, the OpenBoot PROM cannot reliably probe beyond four levels of devices.
- Do not plug a bus-powered hub into another bus-powered hub in a cascading style. A bus-powered hub does not have its own power supply.
- Do not connect a device that requires a large amount of power to a bus-powered hub. These devices might be denied connection to bus-powered hubs or might drain the hub of power for other devices. An example of such a device is a USB diskette device.

SPARC: USB Power Management

Suspending and resuming USB devices is fully supported on SPARC systems. However, do not suspend a device that is busy and never remove a device when the system is powered off under a suspend shutdown.

The USB framework makes a best effort to power manage all devices on SPARC based systems with power management enabled. Power managing a USB device means that the hub driver suspends the port to which the device is connected. Devices that support *remote wake up* can notify the system to wake up everything in the device's path so that the device can be used. The host system could also wake up the device if an application sends an I/O to the device.

All HID devices (keyboard, mouse, hub, and storage devices), hub devices, and storage devices are power managed by default if they support remote wake-up capability. A USB printer is power managed only between two print jobs. Devices that are managed by the generic USB driver (UGEN) are power managed only when they are closed.

When power management is running to reduce power consumption, USB leaf devices are powered down first. After all devices that are connected to a hub's ports are powered down, the hub is powered down after some delay. To achieve the most efficient power management, do not cascade many hubs.

For information about using the SUSPEND/SHUTDOWN key on SPARC systems, see [“USB Keyboards and Mouse Devices” on page 102](#).

Guidelines for Connecting USB Cables

Keep the following guidelines in mind when connecting USB cables:

- USB 3.0 specification does not specify a maximum cable length . However, a USB 3.0 cable is distinctly different from a USB 2.0 cable. It is estimated that shorter cables achieve better performance and three meter cables are recommended to achieve multi-gigabit transfer rates.
- For USB 2.0 devices, always use compliant, fully rated (480 Mbit/sec) 20/28 AWG cables for connecting these devices.
- The maximum cable length that is supported is 5 meters.
- Do not use cable extenders. For best results, use a self-powered hub to extend cable length.

For more information, go to <http://www.usb.org/about/faq>.

Overview of USB Devices

The following overview information is provided in this section:

- “Commonly Used USB Acronyms” on page 105
- “Oracle Solaris USB Architecture” on page 106
- “USB Bus Description” on page 107
- “USB Devices and Driver Classes” on page 108

Universal Serial Bus (USB) was developed by the PC industry to provide a low-cost solution for attaching peripheral devices, such as keyboards, mouse devices, and printers, to a system.

USB connectors are designed to fit only one type of cable, in one way. The primary design motivation for USB was to alleviate the need for multiple connector types for different devices. This design reduces the clutter on the back panel of a system.

Devices connect to USB ports on external USB hubs, or on a root hub that is located on the computer itself. Since hubs have several ports, several branches of a device tree can stem from a hub.

For more information, see [usba\(7D\)](#) or go to <http://www.usb.org/home>.

Commonly Used USB Acronyms

The following table describes the USB acronyms that are used in the Oracle Solaris OS. For a complete description of USB components and acronyms, go to <http://www.usb.org/home>.

Acronym	Definition	For More Information
UGEN	USB generic driver	ugen(7D)
USB	Universal Serial Bus	usb(7D)
USBA	Universal Serial Bus Architecture (Solaris)	usba(7D)
USBAI	USBA Client Driver Interface (Solaris)	N/A
HCD	USB host controller driver	N/A
EHCI	Enhanced Host Controller Interface	ehci(7D)
OHCI	Open Host Controller Interface	ohci(7D)
UHCI	Universal Host Controller Interface	uhci(7D)
XHCI	Extensible Host Controller Interface	xhci(7D)

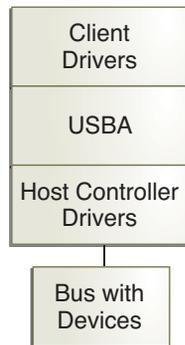
Oracle Solaris USB Architecture

USB devices can be represented as two levels of device tree nodes. A device node represents the entire USB *device*. One or more child *interface* nodes represent the individual USB interfaces on the device.

Driver binding is achieved by using the compatible name properties. For more information, refer to 3.2.2.1 of the IEEE 1275 USB binding and *Writing Device Drivers*. A driver can either bind to the entire device and control all the interfaces, or can bind to just one interface. If no vendor or class driver claims the entire device, a generic USB multi-interface driver is bound to the device-level node. This driver attempts to bind drivers to each interface by using compatible names properties, as defined in section 3.3.2.1 of the IEEE 1275 binding specification.

The Oracle Solaris USB Architecture (USBA) adheres to the USB 1.1, USB 2.0, and USB 3.0 specifications and is part of the Oracle Solaris Device Driver Interface (DDI). The USBA model is similar to Oracle Common SCSI Architecture (SCSA). As the following figure shows, the USBA is a thin layer that provides a generic USB transport-layer abstraction to client drivers, providing them with services that implement core generic USB functionality.

FIGURE 5-1 Oracle Solaris USB Architecture (USBA)



USB Bus Description

The USB specification is openly available and free of royalties. The specification defines the electrical and mechanical interfaces of the bus and the connectors.

USB employs a topology in which hubs provide attachment points for USB devices. The host controller contains the root hub, which is the origin of all USB ports in the system. For more information about hubs, see [“USB Host Controller and Hubs”](#) on page 103.

FIGURE 5-2 USB Physical Device Hierarchy

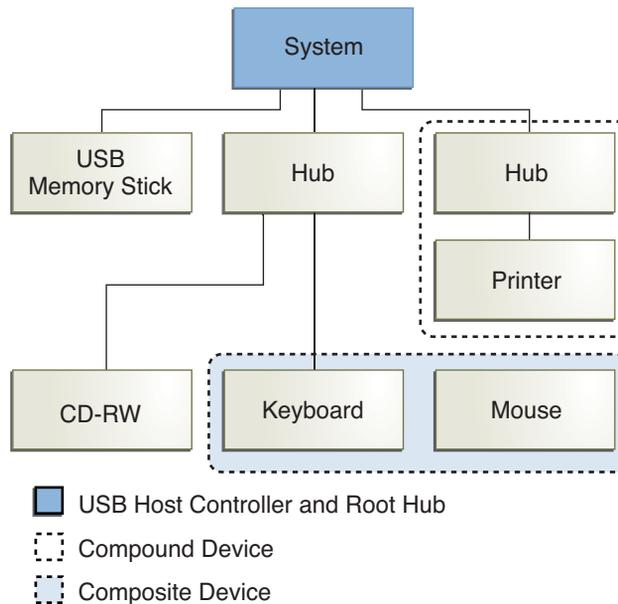


Figure 5–2 shows a system with three active USB ports. The first USB port connects a USB memory stick. The second USB port connects an external hub, which in turn, connects a cdrw device and a composite keyboard/mouse device. As a *composite device*, this keyboard contains a USB controller, which operates both the keyboard and an attached mouse. The keyboard and the mouse share a common USB bus address because they are directed by the same USB controller.

Figure 5–2 also shows an example of a hub and a printer as a *compound device*. The hub is an external hub that is enclosed in the same casing as the printer. The printer is permanently connected to the hub. The hub and printer have separate USB bus addresses.

The device tree path name for some of the devices that are displayed in Figure 5–2 are listed here.

Memory stick	/pci@1f,4000/usb@5/storage@1
Keyboard	/pci@1f,4000/usb@5/hub@2/device@1/keyboard@0
Mouse	/pci@1f,4000/usb@5/hub@2/device@1/mouse@1
cdrw device	/pci@1f,4000/usb@5/hub@2/storage@3
Printer	/pci@1f,4000/usb@5/hub@3/printer@1

USB Devices and Driver Classes

USB devices with similar attributes and services are grouped into device classes. Each device class has a corresponding driver. Devices within a class are managed by the same device driver pair. However, the USB specification also allows for vendor-specific devices that are not part of a specific class.

The HID class contains devices that are user-controlled, such as the following devices:

- Keyboards
- Mouse devices
- Joysticks

The Communication Device class includes the following devices:

- Modems
- Ethernet adapters

Other device classes include the following classes:

- Audio
- Monitor
- Printer
- Storage Device

Each USB device contains descriptors that reflect the class of the device. A device class specifies how its members should behave in configuration and data transfer. You can obtain additional class information by going to <http://www.usb.org/home/>.

For more information about USB devices supported in the Oracle Solaris release, see [usb\(7D\)](#).

Managing USB Mass Storage Devices

The following tasks are described in this section:

- “Using USB Diskette Devices” on page 110
- “Hot-Plugging USB Mass Storage Devices” on page 111
- “Preparing to Use a USB Mass Storage Device” on page 114
- “How to Display USB Device Information” on page 115
- “How to Create a File System on a USB Mass Storage Device ” on page 116
- “How to Modify Partitions and Create a PCFS File System on a USB Mass Storage Device” on page 118
- “How to Create a Solaris Partition and Modify the Slices on a USB Mass Storage Device” on page 121
- “How to Mount or Unmount a USB Mass Storage Device” on page 123
- “Troubleshooting Tips for USB Mass Storage Devices” on page 124
- “Disabling Specific USB Drivers” on page 124
- “How to Disable Specific USB Drivers” on page 125
- “How to Remove Unused USB Device Links” on page 125

The following USB removable mass storage devices are supported:

- CD-RWs
- Hard disks
- DVDs
- Digital cameras
- Diskette devices
- SmartMedia and CompactFlash devices

In previous Oracle Solaris releases, all USB storage devices were identified as removable media devices, which provides many of the following advantages, including automatic mounting. In the Oracle Solaris 11 release, USB mass storage devices are identified as hot-pluggable devices but also enjoy the advantages of USB removable devices that are described in the following sections. For more information about the hot-pluggable behavior, see “[Hot-Plugging USB Mass Storage Devices](#)” on page 111.

Guidelines for managing USB mass storage device are as follows:

- A hot-pluggable device is automatically mounted.
- USB storage devices with standard MS-DOS or Windows (FAT) file systems are supported.
- You can use the user-friendly `rmformat` command to create slices. You can also use the `fdisk` command to partition a USB device, but never use the `format` utility or the `rmformat -F` command to physically format a USB drive.
- Use the `rmformat` command to display all USB devices with media inserted. For example, see “[How to Display USB Device Information](#)” on page 115.
- Non-root users can now access USB storage devices, since the `mount` command is no longer needed. The device is automatically mounted and is available under the `/media` directory.
- These devices can be managed with or without removable media services.
- Disks with FAT file systems can be mounted and accessed. For example:

```
mount -F pcfs /dev/dsk/c2t0d0s0:c /mnt
```
- All USB storage devices are now power managed, except for those that support LOG SENSE pages. Devices with LOG SENSE pages are usually SCSI drives connected through a USB-to-SCSI bridge device.
- Applications might work differently with USB mass storage devices. Keep the following issues in mind when using applications with USB storage devices:
 - Applications might make incorrect assumptions about the size of the media since only smaller devices like diskettes were removable previously.
 - Requests by applications to eject media on devices where this would be inapplicable, such as a hard drive, will succeed and do nothing.
 - If you prefer the behavior in previous Oracle Solaris releases where all USB mass storage were treated as removable media devices, then you can force the old behavior by updating the `/etc/driver/drv/scsa2usb.conf` file.

For more information on using USB mass storage devices, see [scsa2usb\(7D\)](#).

Using USB Diskette Devices

USB diskette devices appear as removable media devices. USB diskette devices are not managed by the `fd` (floppy) driver. Applications that issue `ioctl(2)` calls intended for the `fd` (native floppy) driver will fail. Applications that issue only `read(2)` and `write(2)` calls will succeed. Other applications, such as SunPCI and `rmformat`, also succeed.

The USB diskette device is identified as a SCSI removable media device. The device is available for access under the `/media` directory.

For more information on how to use USB diskette devices, see “[Overview of USB Devices](#)” on page 105.

Hot-Plugging USB Mass Storage Devices

Hot-plugging a device means the device is added or removed without shutting down the operating system or powering off the system. All USB devices are hot-pluggable.

The *hot-pluggable* device attribute identifies those devices that can be connected or disconnected without rebooting the system and configured or unconfigured automatically without user intervention. All USB devices are identified as hot-pluggable devices to gain those benefits. In addition, non-removable media USB devices are no longer identified as removable-media devices and no longer have a *removable-media* attribute.

Non-removable USB storage devices are identified as hot-pluggable devices at the driver level. This behavior means that these devices can be connected or disconnected without rebooting the system and configured or unconfigured automatically without intervention. These changes are made at the kernel level and do not impact the use of these devices. For example, the responsibility of mounting and unmounting these devices is controlled by the removable media management services.

For more information about using these devices, see [scsa2usb\(7D\)](#).

The removable media manager is now aware of hot-plugged devices. You can just plug in the device, which is mounted in a few seconds. If nothing happens, check to see if it is mounted.

Make sure that removable media services are running.

```
# svcs hal dbus rmvolmgr
STATE      STIME      FMRI
online     May_03     svc:/system/dbus:default
online     May_03     svc:/system/hal:default
online     May_03     svc:/system/filesystem/rmvolmgr:default
```

The file system can be mounted from the device if it is valid and it is recognized.

If the file system on the device is not automatically mounted, try a manual mount.

Before hot-removing the device, find the name of the device in the `eject -l` command's alias name. Then eject the device's media. If you don't do this, the device is released and the port is usable again, but the file system on the device might have been damaged.

When you hot-plug a USB device, the device is immediately seen in the system's device hierarchy, as displayed in the `prtconf` command output. When you remove a USB device, the device is removed from the system's device hierarchy, unless you are using the device.

If you are using a device when it is unplugged, the device node remains, but the driver controlling this device stops all activity on the device. Any new I/O activity issued to this device returns an error.

In this situation, the system prompts you to plug in the original device. If the device is no longer available, stop the applications. After a few seconds, the port becomes available again.

Note – Data integrity might be impaired if you remove an active or open device. Always close the device before removing, except the attached keyboard and mouse, which can be moved while active.

▼ How to Add a USB Mass Storage Device

- 1 **Become an administrator.**
- 2 **Connect the USB mass storage device.**
- 3 **Verify that the USB device has been added.**

For example:

```
$ rmformat
Looking for devices...
  1. Logical Node: /dev/rdisk/c3t0d0p0
     Physical Node: /pci@0,0/pci108e,534a@2,1/storage@3/disk@0,0
     Connected Device: SanDisk  Cruzer Micro    0.3
     Device Type: Removable
     Bus: USB
     Size: 245.0 MB
     Label: <None>
     Access permissions: Medium is not write protected.
```

- 4 **Verify that the device is automatically mounted under the /media directory.**

For example:

```
$ ls /media/NONAME
aa  bb
```

You can also use the `rmmount -l` command to list the paths and nicknames of mountable devices. If the device has been mounted under the /media directory, you will see output similar to the following:

```
$ rmmount -l
/dev/dsk/c3t0d0p0:1  rmdisk0,NONAME,/media/NONAME
```

▼ How to Add a USB Camera

If the camera's media uses a PCFS file system, it will be automatically mounted. If the device does not bind to the `scsa2usb` driver, use `libusb` applications for transferring the pictures. For more information, refer to the `/usr/share/doc/libusb/libusb.txt` file.

1 Become an administrator.

2 Plug in and turn on the USB camera.

The system creates a logical device for the camera. After the camera is plugged in, output is written to the `/var/adm/messages` file to acknowledge the device's connection. The system treats the camera as a storage device.

3 Examine the output that is written to the `/var/adm/messages` file.

```
# more /var/adm/messages
```

Examining this output enables you to determine which logical device was created so that you can then use that device to access your images. The output looks similar to the following:

```
Jul 15 09:53:35 buffy usba: [ID 349649 kern.info] OLYMPUS, C-3040ZOOM,
000153719068
Jul 15 09:53:35 buffy genunix: [ID 936769 kern.info] scsa2usb1 is
/pci@0,0/pci925,1234@7,2/storage@2
Jul 15 09:53:36 buffy scsi: [ID 193665 kern.info] sd3 at scsa2usb1:
target 0 lun 0
```

Match the device with a mountable `/dev/dsk` link entry, by doing the following:

```
# ls -l /dev/dsk/c*0 | grep /pci@0,0/pci925,1234@7,2/storage@2
lrwxrwxrwx 1 root root 58 Jun 14 2010 c3t0d0p0 ->
../../devices/pci@0,0/pci925,1234@7,2/storage@2/disk@0,0:a
```

4 Mount the USB camera file system.

The camera's file system is most likely a PCFS file system. If file system is PCFS, then it should be automatically mounted.

- To manually mount the file system on an x86 system, you would use syntax similar to the following:

```
# mount -F pcfs /dev/dsk/c3t0d0p0:c /mnt
```

- To manually mount the file system on a SPARC system, you would use syntax similar to the following:

```
# mount -F pcfs /dev/dsk/c3t0d0s0:c /mnt
```

For information on mounting file systems, see [“Mounting and Unmounting Oracle Solaris File Systems” on page 294](#).

For information on mounting different PCFS file systems, see [`mount_pcfs\(1M\)`](#).

5 (Optional) Verify that the image files are available.

For example:

```
# ls /mnt/DCIM/100OLYMP/  
P7220001.JPG* P7220003.JPG* P7220005.JPG*  
P7220002.JPG* P7220004.JPG* P7220006.JPG*
```

6 (Optional) View and manipulate the image files created by the USB camera.

For example:

```
# /usr/dt/bin/sdtimage P7220001.JPG &
```

7 Unmount the file system before disconnecting the camera.

For example:

```
# umount /mnt
```

8 (Optional) Turn off and disconnect the camera.**▼ How to Remove a USB Mass Storage Device****1 Become an administrator.****2 Stop any active applications that are using the device.****3 Unmount the device.**

```
$ rmount NONAME
```

Or, use the `umount` command as an administrator. For example:

```
# umount /media/NONAME
```

For more information about unmounting a USB device, see [“How to Mount or Unmount a USB Mass Storage Device” on page 123](#).

4 Remove the device.

Preparing to Use a USB Mass Storage Device

You can access information on removable media with or without using removable media services. For information on accessing information on removable media with GNOME's File Manager, see the GNOME desktop documentation.

After the USB device is formatted, it is usually mounted under the `/media/label` directory.

The device nodes are created under the `/dev/rdisk` directory for character devices and under the `/dev/dsk` directory for block devices. Device links are created when the devices are hot-plugged. For more information, see [scsa2usb\(7D\)](#).

If the device cannot be identified by removable media services, then try to manually mount the device. First, use the `rmformat` command to identify the device path and then use the `mount` command to manually mount the device as an administrator.

If a device is mounted by removable media services then you can unmount it by using the `rmumount` command. If the device is manually mounted, then you would unmount it using the `umount` command as an administrator.

For more information about mounting and unmounting USB devices, see “[How to Mount or Unmount a USB Mass Storage Device](#)” on page 123.

▼ How to Display USB Device Information

● Display information about USB devices.

The `prtconf` output in this example has been truncated to only display USB device information.

```
$ prtconf
usb, instance #0
  hub, instance #2
    device, instance #8
      interface (driver not attached)
      printer (driver not attached)
      mouse, instance #14
    device, instance #9
      keyboard, instance #15
      mouse, instance #16
    storage, instance #7
      disk (driver not attached)
      communications, instance #10
      modem (driver not attached)
      data (driver not attached)
  storage, instance #0
    disk (driver not attached)
  storage, instance #1
    disk (driver not attached)
```

Use the `rmformat` command to display USB storage device information:

```
$ rmformat
Looking for devices...
1. Logical Node: /dev/rdisk/c3t0d0p0
   Physical Node: /pci@0,0/pci108e,534a@2,1/storage@3/disk@0,0
   Connected Device: SanDisk  Cruzer Micro    0.3
   Device Type: Removable
   Bus: USB
```

```
Size: 245.0 MB
Label: <None>
Access permissions: Medium is not write protected.
```

▼ How to Create a File System on a USB Mass Storage Device

A USB diskette must be formatted before you can add a file system to it. All other USB mass storage devices just need a file system before they can be used.

Keep the following key points in mind when formatting a USB device:

- Do not use the `rmformat -F` except on a USB diskette.
- If the default slices are not acceptable, use the `rmformat -s` command to create slices. Use the `fdisk` utility to partition a USB device, if needed. For step-by-step instructions, see:
 - [“How to Modify Partitions and Create a PCFS File System on a USB Mass Storage Device” on page 118](#)
 - [“How to Create a Solaris Partition and Modify the Slices on a USB Mass Storage Device” on page 121](#)
- If the USB device is automatically mounted, you will have to unmount it before you can create a file system on the USB device. Use the `rmmount -l` command to identify the device nickname and then the `rmumount` command to unmount the USB device.

Note – Perform Steps 4-5 only if you need to format a USB diskette.

1 Become an administrator.

2 Add the USB device to your system. For information on hot-plugging USB devices, see:

- [“Hot-Plugging USB Mass Storage Devices” on page 111](#)
- [“Hot-Plugging USB Devices With the `cfgadm` Command” on page 129](#)

3 (Optional) Identify the USB device.

```
# rmformat
Looking for devices...
  1. Logical Node: /dev/rdsk/c2t0d0p0
     Physical Node: /pci@0,0/pci108e,534a@2,1/hub@7/floppy@1/disk@0,0
     Connected Device: MITSUMI USB FDD 1039
     Device Type: Floppy drive
     Bus: USB
     Size: 1.4 MB
     Label: <None>
     Access permissions: Medium is not write protected.
```

In this example, the diskette device is `c2t0d0p0`.

4 Insert a diskette into the diskette drive, if necessary.

5 Format the diskette, if necessary.

```
# rmformat -F long raw-device
```

6 Determine the file system type and make sure the device is unmounted. Then, select one of the following:

For more information about unmounting a USB device, see [“How to Mount or Unmount a USB Mass Storage Device” on page 123](#).

- Create a ZFS pool and file system on a USB stick if you need to transfer data to another system.

```
# zpool create c5t0d0 temp-pool
# zfs create temp-pool/data
```

- Create a PCFS file system.

```
# mkfs -F pcfs -o nofdisk,size=size raw-device
```

Specify the `-size` option in 512-byte blocks.

The following example shows how to create a PCFS file system on a 1.4-MB diskette on a SPARC system:

```
# mkfs -F pcfs /dev/rdisk/c2t0d0p0
Construct a new FAT file system on /dev/rdisk/c2t0d0p0: (y/n)? y
```

The following example shows how to create a PCFS file system on a 1.4-MB diskette on an x86 system:

```
# mkfs -F pcfs /dev/rdisk/c2t0d0s2
Construct a new FAT file system on /dev/rdisk/c2t0d0s2: (y/n)? y
```

The following example shows how to create a PCFS file system on a 100-MB USB memory stick on a SPARC system:

```
# mkfs -F pcfs /dev/rdisk/c5t0d0s2:c
```

The following example shows how to create a PCFS file system on a 100-MB USB memory stick on an x86 system:

```
# mkfs -F pcfs /dev/rdisk/c5t0d0p0:c
```

This command can take several minutes to complete.

- Create a legacy UFS file system.

```
# newfs raw-device
```

Consider using the `newfs -f 4096` option or `newfs -T` option for large USB hard disks.

Note – UFS file system overhead consumes a significant portion of space on a diskette, due to a diskette's limited storage capacity.

See “How to Modify Partitions and Create a PCFS File System on a USB Mass Storage Device” on page 118 and “How to Create a Solaris Partition and Modify the Slices on a USB Mass Storage Device” on page 121 for detailed examples of creating a PCFS file system and modifying slices on a USB mass storage device.

▼ How to Modify Partitions and Create a PCFS File System on a USB Mass Storage Device

The following steps describe how to delete an existing partition, create a new partition, and then create a PCFS file system on the USB device. Make sure you backup any data before you perform this task. Note that this procedure includes specific examples to serve as a guideline only. Information that you provide should apply to your particular system.

1 Become an administrator.

2 Start the fdisk utility.

```
# fdisk /dev/rdisk/c3t0d0p0
```

3 Delete the partition by selecting option 3.

```
Total disk size is 29 cylinders
      Cylinder size is 2048 (512 byte) blocks
```

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	=====	=====	=====
1	Active	Solaris2	1	28	28	97

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Enter Selection: 3

4 Choose the partition number to delete.

```
Total disk size is 29 cylinders
      Cylinder size is 2048 (512 byte) blocks
```

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	=====	=====	=====

```

1      Active  Solaris2      1  28    28    97

```

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Specify the partition number to delete (or enter 0 to exit): 1

Partition deleted.

5 Create a partition.

```

Total disk size is 29 cylinders
  Cylinder size is 2048 (512 byte) blocks

```

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	===	=====	===

WARNING: no partitions are defined!

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Enter Selection: 1

6 Select the FAT32 partition type.

```

Total disk size is 29 cylinders
  Cylinder size is 2048 (512 byte) blocks

```

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	===	=====	===

WARNING: no partitions are defined!

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Select the partition type to create:

```

1=SOLARIS2  2=UNIX      3=PCIX05    4=Other
5=DOS12     6=DOS16     7=DOSEXT   8=DOSBIG
9=DOS16LBA A=x86 Boot  B=Diagnostic C=FAT32
D=FAT32LBA E=DOSEXTLBA F=EFI       0=Exit? c

```

7 Specify the percentage of disk to use for this partition.

Total disk size is 29 cylinders
 Cylinder size is 2048 (512 byte) blocks

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	===	=====	===

WARNING: no partitions are defined!

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Select the partition type to create:

Specify the percentage of disk to use for this partition (or type "c" to specify the size in cylinders). **100**

8 Select whether the new partition should be the active partition or an inactive partition.

Total disk size is 29 cylinders
 Cylinder size is 2048 (512 byte) blocks

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	===	=====	===

WARNING: no partitions are defined!

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Select the partition type to create:

Should this become the active partition? If yes, it will be activated each time the computer is reset or turned on.

Please type "y" or "n". **n**

9 Update the disk configuration and exit.

Total disk size is 29 cylinders
 Cylinder size is 2048 (512 byte) blocks

Partition	Status	Type	Cylinders		Length	%
			Start	End		
=====	=====	=====	=====	===	=====	===
1		Win95 FAT32	1	28	28	97

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition

```

3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)
Enter Selection: 5

```

10 Create the PCFS file system on this partition.

Make sure the device is unmounted before creating the new file system. For more information about unmounting a USB device, see [“How to Mount or Unmount a USB Mass Storage Device” on page 123](#).

```

# mkfs -F pcfs -o fat=32 /dev/rdisk/c3t0d0p0:c
Construct a new FAT file system on /dev/rdisk/c3t0d0p0:c: (y/n)? y

```

▼ How to Create a Solaris Partition and Modify the Slices on a USB Mass Storage Device

The following steps illustrate how to create a Solaris partition and modify the slices. Note that this procedure includes specific examples to serve as a guideline only. Information that you provide should apply to your particular system.

Make sure you back up any data before you perform this task.

1 Assume the root role.

2 Start the fdisk utility.

```

# fdisk /dev/rdisk/c5t0d0s2
No fdisk table exists. The default partition for the disk is:

    a 100% "SOLARIS System" partition

Type "y" to accept the default partition, otherwise type "n" to edit the
partition table.
y

```

3 Display the current slices.

For example:

```

# prtvtoc /dev/rdisk/c5t0d0s2
* /dev/rdisk/c5t0d0s2 partition map
*
* Dimensions:
*   512 bytes/sector
*   63 sectors/track
*   255 tracks/cylinder
*  16065 sectors/cylinder
*   5836 cylinders
*   5836 accessible cylinders
*

```

```

* Flags:
* 1: unmountable
* 10: read-only
*
*
* Partition Tag Flags First Sector Last Sector Mount Directory
* 0 0 00 0 93755340 93755339
* 2 0 00 0 93755340 93755339

```

4 Create a text file with the slice information.

For example:

```

slices: 0 = 0, 5GB, "wm", "home" :
        1 = 8225280000, 6GB :
        2 = 0, 44GB, "wm", "backup" :
        6 = 16450560000, 15GB

```

Make sure each slice starts on a cylinder boundary. For example, slice 1 starts at 822280000 bytes, which is the cylinder size in bytes multiplied by 1000.

For more information, see the `-s` option description in `rmformat(1)`.

5 Create the slices by including the slice file created above.

For example:

```
# rmformat -s slice_file /dev/rdisk/c5t0d0s2
```

6 View the new slice information.

For example:

```

# prtvtoc /dev/rdisk/c5t0d0s2
* /dev/rdisk/c5t0d0s2 partition map
*
* Dimensions:
* 512 bytes/sector
* 63 sectors/track
* 255 tracks/cylinder
* 16065 sectors/cylinder
* 5836 cylinders
* 5836 accessible cylinders
*
* Flags:
* 1: unmountable
* 10: read-only
*
* Unallocated space:
* First Sector Last
* Sector Count Sector
* 10485760 5579240 16064999
* 28647912 3482088 32129999
* 63587280 30168060 93755339
*
*
* Partition Tag Flags First Sector Last Sector Mount Directory
* 0 8 00 0 10485760 10485759
* 1 3 01 16065000 12582912 28647911

```

```

2      5      00          0  92274688  92274687
6      4      00    32130000  31457280  63587279

```

▼ How to Mount or Unmount a USB Mass Storage Device

The following steps illustrate how to mount and unmount a USB mass storage device.

1 Become an administrator or the console user.

2 (Optional) Identify the device.

For example:

```

$ rmformat
Looking for devices...
  1. Logical Node: /dev/rdisk/c3t0d0p0
     Physical Node: /pci@0,0/pci108e,534a@2,1/storage@3/disk@0,0
     Connected Device: SanDisk  Cruzer Micro    0.3
     Device Type: Removable
     Bus: USB
     Size: 245.0 MB
     Label: <None>
     Access permissions: Medium is not write protected.

```

In this example, the physical diskette device is `c2t0d0p0`.

3 Select one of the following to mount or unmount a USB mass storage device:

- Mount a USB mass storage device as a console user.

You can use the `rmmount` command with device nicknames, mount points, or device paths, similar to the following:

```

$ rmmount rmdisk0
$ rmmount NONAME
$ rmmount /dev/dsk/c3t0d0p0:1

```

For example:

```

$ rmmount NONAME
NONAME /dev/dsk/c2t0d0p0 mounted
$ ls /media/NONAME
AA.TXT

```

- Unmount a USB mass storage device as a console user.

For example:

```

$ rmmount NONAME
NONAME /dev/dsk/c2t0d0p0 unmounted

```

- Mount a USB mass storage device as superuser.

This example shows how to mount a device with a UFS file system:

```

$ mount /dev/dsk/c1t0d0s2 /mnt

```

This example shows how to mount a device with a PCFS file system on a SPARC system:

```
$ mount -F pcfs /dev/dsk/c3t0d0s2:c /mnt
```

This example shows how to mount a device with a PCFS file system on an x86 system:

```
$ mount -F pcfs /dev/dsk/c3t0d0p0:c /mnt
```

This example shows how to mount a CD with a read-only HSFS file system:

```
$ mount -F hsfs -o ro /dev/dsk/c1t0d0s2 /mnt
```

- Unmount a USB mass storage device as superuser.

First, be sure no one is using the file system on the device.

For example:

```
$ fuser -c -u /mnt
$ umount /mnt
```

4 Eject the device, which is optional for DVD, CD, or diskette devices.

For example:

```
$ eject /dev/rdisk/c1t0d0s2
```

Troubleshooting Tips for USB Mass Storage Devices

Keep the following tips in mind if you have problems adding or removing a USB mass storage device.

Check the `/var/adm/messages` file for failures to enumerate the device. For enumeration failures, possibly, insert the USB hub or remove a hub and connect it directly to a root USB hub.

- If you have problems accessing a device that was connected while the system is running, try the following command:

```
# devfsadm
```
- Do not move devices around if the system has been powered down by a suspend operation. For more information, see [“SPARC: USB Power Management” on page 104](#).
- If a device has been hot removed while in use by applications and is no longer available, then stop the applications. Use the `prtconf` command to see whether the device node has been removed.

Disabling Specific USB Drivers

You can disable specific types of USB devices by disabling their client driver. For example, USB printers can be disabled by disabling the `usbprn` driver that directs them. Disabling `usbprn` does not affect other kinds of devices, such as USB storage devices.

The following table identifies some USB device types and their corresponding drivers.

Device Type	Driver to Disable
Audio	usb_ac and usb_as
HID (usually keyboard and mouse)	hid
Storage	scsa2usb
Printer	usbprn
Serial	usbser_edge

If you disable a driver for a USB device that is still connected to the system, you see a console message similar to the following:

```
usb10: WARNING: usba: no driver found for device name
```

▼ How to Disable Specific USB Drivers

- 1 **Become an administrator.**
- 2 **Exclude the driver alias entry from the `/etc/system` file.**

For example, include the following exclude statement to exclude the `usbprn` driver:

```
exclude: usbprn
```

- 3 **Reboot the system.**

```
# init 6
```

▼ How to Remove Unused USB Device Links

Use this procedure if a USB device is removed while the system is powered off. Removing the USB device while the system is powered off can leave device links for devices that do not exist.

- 1 **Become an administrator.**
- 2 **Close all applications that might be accessing the device.**
- 3 **Remove the unused links for a specific USB class.**

For example:

```
# devfsadm -C -c audio
```

Or, just remove the dangling links:

```
# devfsadm -C
```

Using USB Audio Devices

The following tasks are provided in this section:

- “Hot-Plugging Multiple USB Audio Devices” on page 127
- “How to Add USB Audio Devices” on page 127
- “How to Identify Your System's Primary Audio Device” on page 128
- “How to Change the Primary USB Audio Device” on page 128
- “Troubleshooting USB Audio Device Problems” on page 129

Note – USB 3.0 does *not* include support for audio devices.

For information about USB audio support in specific Oracle Solaris releases, see “[About USB Support in Oracle Solaris](#)” on page 100.

Oracle Solaris USB audio support is implemented by a pair of cooperating drivers, `usb_ac` and `usb_as`. The audio control driver, `usb_ac`, is a Solaris USB Architecture compliant client driver that provides the controlling interface to user applications. The audio streaming driver, `usb_as`, processes audio data messages during play and record. It sets sample frequency and precision, and encodes requests from the `usb_ac` driver. Both drivers comply with the USB audio class 1.0 specification.

Some audio devices can set volume under software control. A STREAMS module, `usb_ah`, is pushed on top of the HID driver for managing this function.

Oracle Solaris supports USB audio devices that are play-only, record-only, or record and play.

Hot-plugging of USB audio devices is supported, as follows:

- For fully supported audio data format information, see [usb_ac\(7D\)](#).

The primary audio device is `/dev/audio`. You can verify that `/dev/audio` is pointing to USB audio by using the following command:

```
root% mixerctl
Device /dev/audioctl:
  Name    = USB Audio
  Version = 1.0
  Config  = external
```

```
Audio mixer for /dev/audioctl is enabled
```

After you connect your USB audio devices, you access them with the `audioplay` and `audiorecord` command through the `/dev/sound/N` device links.

Note that the `/dev/audio` and `/dev/sound/N` devices can refer to speakers, microphones, or combination devices. If you refer to the incorrect device type, the command fails. For example, the `audioplay` command fails if you try to use it with a microphone.

You can select a specific default audio device for most Oracle audio applications, such as `audioplay` and `audiorecord`, by setting the `AUDIODEV` shell variable or by specifying the `-d` option for these commands. However, setting `AUDIODEV` does not work for third-party applications that have `/dev/audio` hardcoded as the audio file.

When you plug in a USB audio device, it automatically becomes the primary audio device, `/dev/audio`, unless `/dev/audio` is in use. For instructions on changing `/dev/audio` from on-board audio to USB audio and vice versa, refer to “[How to Change the Primary USB Audio Device](#)” on page 128, and [usb_ac\(7D\)](#).

Hot-Plugging Multiple USB Audio Devices

If a USB audio device is plugged into a system, it becomes the primary audio device, `/dev/audio`. It remains the primary audio device even after the system is rebooted. If additional USB audio devices are plugged in, the last one becomes the primary audio device.

For additional information on troubleshooting USB audio device problems, see [usb_ac\(7D\)](#).

▼ How to Add USB Audio Devices

1 Plug in the USB speaker.

The primary audio device, `/dev/audio`, points to the USB speaker.

```
$ ls -l /dev/audio
lrwxrwxrwx 1 root  root  10 Feb 13 08:46 /dev/audio -> usb/audio0
```

2 (Optional) Remove the speaker. Then, plug it back in.

If you remove the speaker, the `/dev/audio` device reverts back to on-board audio.

```
$ ls -l /dev/audio
lrwxrwxrwx 1 root  root   7 Feb 13 08:47 /dev/audio -> sound/0
```

3 Add a USB microphone.

```
$ ls -l /dev/audio
lrwxrwxrwx 1 root  root  10 Feb 13 08:54 /dev/audio -> usb/audio1
```

▼ How to Identify Your System's Primary Audio Device

This procedure assumes that you have already connected the USB audio devices.

● Examine your system's new audio links.

- Display your system's new audio links with the `ls` command.

For example:

```
$ ls -lt /dev/audio*
lrwxrwxrwx 1 root root    7 Jul 23 15:46 /dev/audio -> usb/audio0
lrwxrwxrwx 1 root root   10 Jul 23 15:46 /dev/audiocctl ->
usb/audiocctl0/
% ls -lt /dev/sound/*
lrwxrwxrwx 1 root root    74 Jul 23 15:46 /dev/sound/1 ->
.././devices/pci@1f,4000/usb@5/hub@1/device@3/sound-control@0:...
lrwxrwxrwx 1 root root    77 Jul 23 15:46 /dev/sound/1ctl ->
.././devices/pci@1f,4000/usb@5/hub@1/device@3/sound-control@0:...
lrwxrwxrwx 1 root other   66 Jul 23 14:21 /dev/sound/0 ->
.././devices/pci@1f,4000/ebus@1/SUNW,CS4231@14,200000:sound,audio
lrwxrwxrwx 1 root other   69 Jul 23 14:21 /dev/sound/0ctl ->
.././devices/pci@1f,4000/ebus@1/SUNW,CS4231@14,200000:sound,audiocctl
$
```

Notice that the primary audio device, `/dev/audio`, is pointing to the newly plugged in USB audio device, `/dev/usb/audio0`.

- You can also examine your system's USB audio devices with the `prtconf` command and look for the USB device information.

```
$ prtconf
.
.
.
usb, instance #0
  hub, instance #0
    mouse, instance #0
    keyboard, instance #1
    device, instance #0
      sound-control, instance #0
      sound, instance #0
      input, instance #0
.
.
.
```

▼ How to Change the Primary USB Audio Device

● Select one of the following to change the primary USB audio device:

- If you want the on-board audio device to become the primary audio device, remove the USB audio devices. The `/dev/audio` link then points to the `/dev/sound/0` entry. If the

`/dev/sound/0` entry is not the primary audio device, then either shut down the system and use the `boot -r` command, or run the `devfsadm -i` command as root.

- If you want the USB audio device to become primary audio device, just plug it in and check the device links.

Troubleshooting USB Audio Device Problems

Sometimes, USB speakers do not produce any sound, even though the driver is attached and the volume is set to high. Hot-plugging the device might not change this behavior.

The workaround is to power cycle the USB speakers.

Key Points of Audio Device Ownership

Keep the following key points of audio device ownership in mind when working with audio devices:

- When you plug in a USB audio device and you are logged in to the console, the console is the owner of the `/dev/*` entries. This situation means you can use the audio device, as long as you are logged in to the console.
- If you are not logged in to the console when you plug in a USB audio device, root becomes the owner of the device. However, if you log in to the console and attempt to access the USB audio device, then device ownership changes to the console. For more information, see [logindevperm\(4\)](#).
- When you remotely log in with the `rlogin` command and attempt to access the USB audio device, the ownership does not change. This situation means that, for example, unauthorized users cannot listen to conversations over a microphone that is owned by someone else.

Hot-Plugging USB Devices With the `cfgadm` Command

The following tasks are provided in this section:

- “How to Display USB Bus Information (`cfgadm`)” on page 130
- “How to Unconfigure a USB Device” on page 131
- “How to Configure a USB Device” on page 132
- “How to Logically Disconnect a USB Device” on page 132
- “How to Logically Connect a USB Device” on page 133
- “How to Logically Disconnect a USB Device Subtree” on page 133
- “How to Reset a USB Device” on page 134
- “How to Change the Default Configuration of a Multi-Configuration USB Device” on page 134

You can add and remove a USB device from a running system without using the `cfgadm` command. However, a USB device can also be *logically* hot-plugged without physically removing the device. This scenario is convenient when you are working remotely and you need to disable or reset a non functioning USB device. The `cfgadm` command also provides a way to display the USB device tree, including manufacturer and product information.

The `cfgadm` command displays information about *attachment points*, which are locations in the system where dynamic reconfiguration operations can occur.

An attachment point consists of the following:

- An occupant, which represents a hardware resource, such as a USB device, that might be configured into the system
- A receptacle, which is the location that accepts the occupant, such as a USB port

Attachment points are represented by logical and physical attachment point IDs (`Ap_Ids`). The physical `Ap_Id` is the physical path name of the attachment point. The logical `Ap_Id` is a user-friendly alternative for the physical `Ap_Id`. For more information on `Ap_Ids`, see [`cfgadm_usb\(1M\)`](#).

The `cfgadm` command provides the following USB device status information.

Receptacle State	Description
empty/unconfigured	The device is not physically connected.
disconnected/unconfigured	The device is logically disconnected and unavailable, even though the device could still be physically connected.
connected/unconfigured	The device is logically connected, but unavailable. The device is visible in <code>prtconf</code> output.
connected/configured	The device is connected and available.

The following sections describe how to hot-plug a USB device through the software with the `cfgadm` command. All of the sample USB device information in these sections has been truncated to focus on relevant information.

▼ How to Display USB Bus Information (`cfgadm`)

For examples of using the `prtconf` command to display USB configuration information, see [“How to Display USB Device Information”](#) on page 115.

1 Display USB bus information.

For example:

```
$ cfgadm
Ap_Id           Type           Receptacle   Occupant   Condition
usb0/4.5       usb-hub        connected    configured ok
usb0/4.5.1     usb-device    connected    configured ok
usb0/4.5.2     usb-printer   connected    configured ok
usb0/4.5.3     usb-mouse     connected    configured ok
usb0/4.5.4     usb-device    connected    configured ok
usb0/4.5.5     usb-storage   connected    configured ok
usb0/4.5.6     usb-communi   connected    configured ok
usb0/4.5.7     unknown      empty        unconfigured ok
```

In the preceding example, `usb0/4.5.1` identifies a device connected to port 1 of the second-level external hub, which is connected to port 5 of first-level external hub, which is connected to the first USB controller's root hub, port 4.

2 Display specific USB device information.

For example:

```
$ cfgadm -l -s "cols=ap_id:info"
Ap_Id           Information
usb0/4.5.1     Mfg: Inside Out Networks Product: Edgeport/421 NConfigs: 1
Config: 0 : ...
usb0/4.5.2     Mfg: <undef> Product: <undef> NConfigs: 1 Config: 0 ...
usb0/4.5.3     Mfg: Mitsumi Product: Apple USB Mouse NConfigs: 1
Config: 0 ...
usb0/4.5.4     Mfg: NMB Product: NMB USB KB/PS2 M NConfigs: 1 Config: 0
usb0/4.5.5     Mfg: Hagiwara Sys-Com Product: SmartMedia R/W NConfigs: 1
Config: 0 : ...
usb0/4.5.6     Mfg: 3Com Inc. Product: U.S.Robotics 56000 Voice USB Modem
NConfigs: 2 ...
usb0/4.5.7
```

▼ How to Unconfigure a USB Device

You can unconfigure a USB device that is still physically connected to the system. However, a driver will never attach to the device. Note that a USB device remains in the `prtconf` output even after that device is unconfigured.

1 Become an administrator.

2 Unconfigure the USB device.

For example:

```
# cfgadm -c unconfigure usb0/4.7
Unconfigure the device: /devices/pci@8,700000/usb@5,3/hub@4:4.7
This operation will suspend activity on the USB device
Continue (yes/no)? y
```

3 Verify that the device is unconfigured.

For example:

```
# cfgadm
Ap_Id                Type      Receptacle  Occupant  Condition
usb0/4.5             usb-hub   connected   configured ok
usb0/4.5.1           usb-device connected   configured ok
usb0/4.5.2           usb-printer connected   configured ok
usb0/4.5.3           usb-mouse  connected   configured ok
usb0/4.5.4           usb-device  connected   configured ok
usb0/4.5.5           usb-storage connected   configured ok
usb0/4.5.6           usb-communi connected   configured ok
usb0/4.5.7           unknown   empty       unconfigured ok
usb0/4.6             usb-storage connected   configured ok
usb0/4.7             usb-storage connected   unconfigured ok
```

▼ How to Configure a USB Device

1 Become an administrator.

2 Configure a USB device.

For example:

```
# cfgadm -c configure usb0/4.7
```

3 Verify that the USB device is configured.

For example:

```
# cfgadm usb0/4.7
Ap_Id                Type      Receptacle  Occupant  Condition
usb0/4.7             usb-storage connected   configured ok
```

▼ How to Logically Disconnect a USB Device

If you want to remove a USB device from the system and the `prtconf` output, but you are not physically near the system, just logically disconnect the USB device. The device is still physically connected. However, the device is logically disconnected, unusable, and not visible to the system.

1 Become an administrator.

2 Disconnect a USB device.

For example:

```
# cfgadm -c disconnect -y usb0/4.7
```

3 Verify that the device is disconnected.

For example:

```
# cfgadm usb0/4.7
Ap_Id          Type      Receptacle  Occupant    Condition
usb0/4.7      unknown  disconnected  unconfigured ok
```

▼ How to Logically Connect a USB Device

Use this procedure to logically connect a USB device that was previously logically disconnected or unconfigured.

1 Become an administrator.**2 Connect a USB device.**

For example:

```
# cfgadm -c configure usb0/4.7
```

3 Verify that the device is connected.

For example:

```
# cfgadm usb0/4.7
Ap_Id          Type      Receptacle  Occupant    Condition
usb0/4.7      usb-storage  connected   configured  ok
```

The device is now available and visible to the system.

▼ How to Logically Disconnect a USB Device Subtree

Use this procedure to disconnect a USB device subtree, which is the hierarchy (or tree) of devices below a hub.

1 Become an administrator.**2 Remove a USB device subtree.**

For example:

```
# cfgadm -c disconnect -y usb0/4
```

3 Verify that the USB device subtree is disconnected.

For example:

```
# cfgadm usb0/4
Ap_Id          Type      Receptacle  Occupant    Condition
usb0/4        unknown  disconnected  unconfigured ok
```

▼ How to Reset a USB Device

If a USB device behaves erratically, use the `cfgadm` command to reset the device, which logically removes and recreates the device.

- 1 **Become an administrator.**
- 2 **Make sure that the device is not in use.**
- 3 **Reset the device.**

For example:

```
# cfgadm -x usb_reset -y usb0/4.7
```

- 4 **Verify that the device is connected.**

For example:

```
# cfgadm usb0/4.7
Ap_Id          Type          Receptacle  Occupant    Condition
usb0/4.7      usb-storage   connected   configured  ok
```

▼ How to Change the Default Configuration of a Multi-Configuration USB Device

Keep the following in mind when working with multi-configuration USB devices:

- A USB device configuration defines how a device presents itself to the operating system. This method is different from system device configurations discussed in other `cfgadm` sections.
- Some USB devices support multiple configurations, but only one configuration can be active at a time.
- Multi-configuration devices can be identified by examining the `cfgadm -lv` output. `Nconfigs` will be greater than 1.
- The default USB configuration is configuration 1. The current configuration is reflected in `cfgadm -lv` output as `Config`.
- Changes to the default configuration persist across reboots, hot-removes, and the reconfiguration of the device, as long as the device is reconnected to the same port.

- 1 **Make sure that the device is not in use.**
- 2 **Change the default USB configuration.**

For example:

```
# cfgadm -x usb_config -o config=2 usb0/4
Setting the device: /devices/pci@1f,0/usb@c,3:4
```

to USB configuration 2
This operation will suspend activity on the USB device
Continue (yes/no)? **yes**

3 Verify that the device changed.

For example:

```
# cfgadm -lv usb0/4
Ap_Id  Receptacle  Occupant    Condition  Information  When  Type
      Busy      Phys_Id
usb0/4  connected   unconfigured ok      Mfg: Sun 2000
Product: USB-B0B0 aka Robotech
With 6 EPPS High Clk Mode  NConfigs: 7  Config: 2  : EVAL Board Setup
unavailable
usb-device  n          /devices/pci@1f,0/usb@c,3:4
```

Note that Config: now shows 2.

Using InfiniBand Devices (Overview/Tasks)

This chapter provides general overview and step-by-step instructions for using InfiniBand (IB) devices in your network.

This is a list of the overview information in this chapter.

- “Overview of InfiniBand Devices” on page 137
- “Dynamically Reconfiguring IB Devices (c f g adm)” on page 141

For information on the procedures associated with using IB devices, see the following:

- “Dynamically Reconfiguring IB Devices (Task Map)” on page 139
- “Using the uDAPL Application Interface With InfiniBand Devices” on page 150
- “Administering IPoIB Devices (d \ adm)” on page 152
- “Monitoring and Troubleshooting IB Devices” on page 155

For general information about dynamic reconfiguration and hot-plugging, see [Chapter 4](#), “Dynamically Configuring Devices (Tasks).”

Overview of InfiniBand Devices

InfiniBand (IB) is a new I/O technology based on switched fabrics. It provides a high bandwidth, low latency interconnect for attaching I/O devices to hosts and for host-to-host communication. IB devices are managed by the Solaris IB nexus driver.

Support for the following devices is provided in Oracle Solaris 11:

- IP over IB (IPoIB) devices – Enables the ability to transport the IP packets over IB connections. This feature is implemented by the `ibp(7D)` driver
- Socket Direct Protocol (SDP) – Provides support for sockets over IB
- Reliable Datagram Service version 3 (RDSv3) and Reliable Datagram Service (RDS)
- NFS over Remote Direct Memory Access (NFSoverRDMA) - Provides NFS services over IB using RDMA

- iSCSI Extension for RDMA (iSER) – Provides the RDMA data transfer capability to the iSCSI protocol
- User Direct Access Programming Language (uDAPL)
- Open Fabric User Verb (OFUV)
- Ethernet over IB (EoIB)

The IB nexus driver queries the Solaris IB Device Manager (IBDM) for services, referred in this guide as *communication services*, to enumerate the IB Port, HCA_SVC, and IB VPPA devices.

The IB partition link represents a new *part* class of data link and this object is managed by using the new `dladm` subcommands. An IB partition link can be created on top of an IB physical link, one per each `P_Key` on the port. The partition links are used for data transfers.

The Port devices bind a communication service to a given `port#` of a Host Channel Adapter (HCA). The VPPA devices bind a communication service to a `port#`, `P_key#` combination instead. The HCA_SVC devices bind a communication service to a given HCA. Note that the Port devices and the HCA_SVC devices always use a `P_key` (partition key) whose value is zero. The Port, HCA_SVC, and VPPA devices are children of the HCA and are enumerated through the `ib.conf` file. For more information, see [ib\(7D\)](#).

The Input Output Controller (IOC) devices are children of the IB nexus driver and are part of an I/O unit. The pseudo devices are also children of the IB nexus driver and refer to all other devices that provide their own configuration files to enumerate. For more information, see [ib\(4\)](#).

The possible IB device tree path name(s) are listed in the following table.

IOC device	<code>/ib/ioc@1730000007F510C,1730000007F50</code>
IB pseudo device	<code>/ib/driver@unit-address</code>
IB VPPA device	<code>/pci@1f,2000/pci@1/pci15b3,5a44@0/ibport@ ,port#,P_key,service</code>
IB HCA_SVC device	<code>/pci@1f,2000/pci@1/pci15bc,5a44@0/ibport@0,0,service</code>
IB Port device	<code>/pci@1f,2000/pci@1/pci15b3,5a44@0/ibport@<port#>,0,service</code>
HCA	<code>/pci@1f,2000/pci@1/pci15b3,5a44@0</code>

Note that the IB HCA_SVC devices have zero as the `port#` and the `P_key`.

The IB components in the preceding table are described as follows:

services Is a communication service. For example, `ipib` is the communication service used by the `ibd` kernel client driver.

<i>P_key</i>	Is the partition link key value being used.
<i>port</i>	Is the port number.
<i>unit-address</i>	Refers to IB kernel client driver's property by this name specified in its <code>driver.conf</code> file. For more information, see driver.conf(4) .

For information about using IB diagnostic commands and utilities, see “[Monitoring and Troubleshooting IB Devices](#)” on page 155.

InfiniBand Software Packages

The IB related software packages are as follows:

- `driver/infiniband/connectx` – Mellanox ConnectX Family InfiniBand HCA and 10GbE NIC drivers
- `system/io/infiniband/ethernet-over-ib` – InfiniBand device driver implementing Ethernet over InfiniBand
- `system/io/infiniband/ib-device-mgt-agent` – InfiniBand Device Manager Agent
- `system/io/infiniband/ib-sockets-direct` – InfiniBand layered Sockets Direct Protocol
- `system/io/infiniband/ip-over-ib` – Network device driver supporting the IP over InfiniBand (IPoIB) protocol
- `system/io/infiniband/open-fabrics` – Open Fabrics kernel components
- `system/io/infiniband/reliable-datagram-sockets-v3` – Reliable Datagram Sockets (RDSv3)
- `system/io/infiniband/reliable-datagram-sockets` – Reliable Datagram Sockets
- `system/io/infiniband/rpc-over-rdma` – InfiniBand RPC over RDMA Driver
- `system/io/infiniband/udapl` – UDAPL library and commands
- `system/io/infiniband` – InfiniBand Framework

Dynamically Reconfiguring IB Devices (Task Map)

Task	Description	For Instructions
Display IB device information.	Display information about the IB devices on your system.	“How to Display IB Device Information” on page 142
Configure or unconfigure a port or VPPA device.	Select one of the following:	

Task	Description	For Instructions
	<p>Unconfigure a port or a VPPA device.</p> <p>Configure a port or a VPPA device.</p>	<p>“How to Unconfigure an IB Port, HCA_SVC, or a VPPA Device” on page 144</p> <p>“How to Configure a IB Port, HCA_SVC, or a VPPA Device” on page 144</p>
Configure or unconfigure an IB pseudo device.	<p>Select one of the following:</p> <p>Unconfigure an IB pseudo device.</p> <p>Configure an IB pseudo device.</p>	<p>“How to Unconfigure an IB Pseudo Device” on page 145</p> <p>“How to Configure an IB Pseudo Device” on page 145</p>
Display kernel IB clients of an HCA.	You might need to display information about kernel IB clients of an HCA, particularly if you're going to unconfigure an HCA.	“How to Display Kernel IB Clients of an HCA” on page 146
Configure or unconfigure an IB HCA or EoIB interface.	<p>Select one of the following:</p> <p>Unconfigure IB devices that are connected to an HCA.</p> <p>Dynamically reconfigure the HCA when EoIB devices active.</p> <p>Unconfigure the not last IB HCA with EoIB interface(s) active.</p> <p>Configure IB devices that are connected to an HCA.</p>	<p>“How to Dynamically Reconfigure an HCA With Active EoIB Devices” on page 146</p> <p>“How to Reconfigure and Restore an EoIB Interface After Hot Removal” on page 147</p> <p>“Configuring an IB HCA” on page 148</p>
Update the IB P_key tables.	If the P_key table information of a HCA port changes, IBTF and IBDM need to be notified so that their internal P_key databases are updated.	“How to Update the IB P_key Tables” on page 148
Display IB communication services	Display the IB communication services that are currently in use by the IBTF.	“How to Display IB Communication Services” on page 149
Add or remove a VPPA communication service.	Select one of the following:	

Task	Description	For Instructions
	Add a VPPA communication service.	“How to Add a VPPA Communication Service” on page 149
	Remove a VPPA communication service.	“How to Remove an Existing IB Port, HCA_SVC, or a VPPA Communication Service” on page 150

Dynamically Reconfiguring IB Devices (c f g a d m)

One can configure or unconfigure an IB device from a running system by using the `c f g a d m` CLI only. This command also provides a way to display the IB fabric, manage communication services, and update `P_key` table databases. For more information, see [c f g a d m_ib\(1M\)](#).

The `c f g a d m` CLI manages dynamic reconfiguration, referred to in this guide as DR, of the entire IB fabric as seen by a host. The `c f g a d m` operations are supported on all the IB devices, such as Port, VPPA, HCA_SVC, IOC, and pseudo devices.

The `c f g a d m` command displays information about attachment points (`Ap_Ids`), which are locations in the system where DR operations can occur. For details on the `Ap_Ids` that `c f g a d m` supports, see `c f g a d m_ib. 1M`. Note that all IB `Ap_Ids` are shown as connected.

The `c f g a d m` command provides the following IB device status information.

Receptacle State	Description
connected/configured/ok	The device is connected and available. The <code>devinfo</code> node is present.
connected/unconfigured/unknown	The device is unavailable and no <code>devinfo</code> node or device driver exists for this device. Or, the device was never configured for use by <code>ib_nexus</code> driver. The device might be known to the IB Device Manager.

All of the procedures in this section require administrative privileges that are not generally granted to user accounts. For more information, see [“How to Use Your Assigned Administrative Rights” in Oracle Solaris 11.1 Administration: Security Services](#).

The following sections describe how to dynamically reconfigure (DR) IB devices with the `c f g a d m` command. All of the sample IB device information in these sections has been truncated to focus on relevant information.

Unconfiguring IB Device Considerations

An actual dynamic reconfiguration (DR) of an HCA is beyond the scope of the IB `cfgadm` plugin. Although DR of an HCA can be achieved by using the plugin of the underlying bus. For example, a PCI based HCA can use the `cfgadm_pci` command. For more information, see [cfgadm_pci\(1M\)](#).

Note – RDSv3 does not support unconfiguring a HCA. If the system has a RDSv3 driver installed at the time of DR, unconfiguring the HCA fails as shown below.

```
# cfgadm -c unconfigure ib::rdsv3,0
This operation will suspend activity on the IB device
Continue (yes/no)? yes
cfgadm: Hardware specific failure: unconfigure operation failed ap_id: /devices/ib:fabric::rdsv3,0

# cfgadm -c unconfigure PCI-EM0
cfgadm: Component system is busy, try again: unconfigure failed
```

Workaround:

On a production system, you must remove the RDSv3 driver before the HCA DR operation and reboot the system.

```
# rem_drv rdsv3
Device busy
Cannot unload module: rdsv3
Will be unloaded upon reboot.
```

```
# init 6
```

▼ How to Display IB Device Information

You can use the `prtconf` command to display general information about IB devices. For example:

```
$ prtconf
.
.
.
  ib, instance #0
    rpcib, instance #0
    rdsib, instance #0
    daplt, instance #0
    rdsv3, instance #0
    sdplib, instance #0
    eibnx, instance #0
    sol_umad, instance #0
    sol_uverbs, instance #0
    iser, instance #0
```

```

.
.
.
pci15b3,673c, instance #0
    ibport, instance #0
    ibport, instance #1

```

In the above example, pci15b3,673c refers to an IB HCA.

Use the following steps to display specific IB device information.

1 Become an administrator.

2 Display IB fabric information.

For example:

```

# cfgadm -al
Ap_Id      Type      Receptacle  Occupant  Condition
ib         IB-Fabric connected   configured ok
hca:21280001A0A478 IB-HCA    connected   configured ok
ib::21280001A0A47A,0,ipib IB-PORT   connected   configured ok
ib::21280001A0A479,0,ipib IB-PORT   connected   configured ok
ib::1730000008070,0,hnfs IB-HCA_SVC connected   configured ok
ib::daplt,0      IB-PSEUDO connected   configured ok
ib::eibnx,0      IB-PSEUDO connected   configured ok
ib::iser,0       IB-PSEUDO connected   configured ok
ib::rdsib,0      IB-PSEUDO connected   configured ok
ib::rdsv3,0      IB-PSEUDO connected   configured ok
ib::rpcib,0      IB-PSEUDO connected   configured ok
ib::sdpib,0      IB-PSEUDO connected   configured ok
ib::sol_umad,0   IB-PSEUDO connected   configured ok
ib::sol_uverbs,0 IB-PSEUDO connected   configured ok

```

In the above example output, the components are described as follows:

Ap_Id ib::21280001A0A47A,0,ipib	Identifies an IB port device that is connected to port GUID and is bound to the ipib service.
Ap_ID ib::sdpib,0	Identifies a pseudo device.
Ap_ID hca:21280001A0A478	Identifies an HCA device.
Ap_ID ib::1730000008070,0,hnfs	Identifies an IB HCA_SVC device that is bound to the hnfs service.
Ap_Id ib::ibgen,0	Identifies a pseudo device.

3 Display specific IB device information.

For example, for an IB port device:

```

# cfgadm -al -s "cols=ap_id:info" ib::21280001A0A47A,0,ipib
Ap_Id      Information
ib::21280001A0A47A,0,ipib ipib

```

For example, for an IB HCA device:

```
# cfgadm -al -s "cols=ap_id:info" hca::1730000008070
Ap_Id                               Information
hca::1730000008070                 VID: 0x15b3, PID: 0x5a44, #ports: 0x2,
port1 GUID: 0x1730000008071, port2 GUID: 0x1730000008072
```

The preceding output displays the number of ports and their GUIDs.

▼ How to Unconfigure an IB Port, HCA_SVC, or a VPPA Device

Use the following steps if you want to remove an IB Port, HCA_SVC, or a VPPA device from the system.

The example below illustrates how to unconfigure an IB Port device, but the same procedure applies to VPPA and HCA_SVC devices as well.

- 1 **Become an administrator.**
- 2 **Unconfigure virtual IB port devices.**

For example:

```
# cfgadm -c unconfigure ib::1730000007F51,*0*,ipib
Unconfigure the device: /devices/ib:fabric::1730000007F51,*0*,ipib
This operation will suspend activity on the IB device
Continue (yes/no)? Y
```

- 3 **Verify that the device is disconnected.**

For example:

```
# cfgadm -a ib::1730000007F51,*0*,ipib
Ap_Id                               Type      Receptacle Occupant    Condition
ib::1730000007F51,*0*,ipib IB-VPPA   connected  unconfigured unknown
```

▼ How to Configure a IB Port, HCA_SVC, or a VPPA Device

Use the following steps if you want to configure an IB Port, HCA_SVC, or a VPPA device on the system.

The example below illustrates how to configure a VPPA device, but similar steps can be used to configure an IB Port and HCA_SVC devices as well.

- 1 **Become an administrator.**

2 Configure the virtual IB port devices.

For example:

```
# cfgadm -c configure ib::1730000007F51,*0*,ipib
```

3 Verify that the device is connected.

For example:

```
# cfgadm -a ib::1730000007F51,*0*,ipib
Ap_Id                Type                Receptacle Occupant  Condition
ib::1730000007F51,*0*,ipib  IB-VPPA            connected  configured ok
```

Note – A `cfgadm` based configure or unconfigure operation of IB Port and HCA_SVC devices is similar to the preceding examples for an IB VPPA device.

▼ How to Unconfigure an IB Pseudo Device

Use the following steps if you want to remove an IB pseudo device from the system.

1 Become an administrator.**2 Unconfigure the IB pseudo device.**

For example:

```
# cfgadm -c unconfigure ib::ibgen,0
Unconfigure the device: /devices/ib:fabric::ibgen,0
This operation will suspend activity on the IB device
Continue (yes/no)? Y
```

3 Verify that the device is disconnected.

```
# cfgadm -a ib::ibgen,0
Ap_Id                Type                Receptacle Occupant  Condition
ib::ibgen,0          IB-PSEUDO           connected  unconfigured unknown
```

▼ How to Configure an IB Pseudo Device

Use the following steps to configure an IB pseudo device.

1 Become an administrator.**2 Configure the IB pseudo device.**

For example:

```
# cfgadm -yc configure ib::ibgen,0
```

3 Verify that the device is connected.

For example:

```
# c f g a d m - a i b : : i b g e n , 0
Ap_Id                Type           Receptacle Occupant   Condition
ib::ibgen,0         IB-PSEUDO     connected   configured ok
```

▼ How to Display Kernel IB Clients of an HCA

The following IB `c f g a d m` plugin command can be invoked to list kernel IB clients using this HCA. Note that the last column would show a “yes” if a kernel IB client uses another HCA. IB Managers and kernel clients that do not use the HCA are shown with an `Ap_Id` of “-”.

● Display kernel IB clients of an HCA.

For example:

```
$ c f g a d m - x l i s t _ c l i e n t s h c a : 1 7 3 0 0 0 0 0 7 F 5 0
Ap_Id                IB Client           Alternate HCA
ib::1730000007F51D0  ibgen               no
ib::1730000007F51D1  ibgen               no
ib::1730000007F51,0,ipib  ibd                 no
ib::ibgen,0          ibgen               no
-                    ibdm                no
-                    ibmf                no
-                    nfs/ib              no
```

▼ How to Dynamically Reconfigure an HCA With Active EoIB Devices

When a dynamic reconfiguration (DR) operation is attempted on a system and a Connect-X family IB HCA is in use with active EoIB interfaces, (for example, EoIB datalinks are plumbed or VNICs are created over the datalinks on the HCA), the DR operation succeeds if either of the following conditions is true:

- This is the only Connect-X family HCA on the system
- Multiple Connect-X HCAs exist on the system, but this adapter is the last one to be unconfigured

Under any other condition, a dynamic reconfiguration (DR) operation on an IB HCA with active EoIB interfaces fails with a `c f g a d m` error message.

If the DR operation fails, you must unplumb the active EoIB interfaces and delete any VNICs on the datalink and retry the DR operation. In addition, in such a case, if a Connect-X family HCA is replaced in the same slot after the DR operation and configured again using the `c f g a d m` or

hotplug commands, you must replumb the EoIB datalinks that were unplumbed earlier and recreate any VNICs that were previously deleted.

Note that a functioning EoIB interface requires that the HCA port GUID be listed in the appropriate config file on the Sun Network QDR Infiniband Gateway Switch. So, after a hot-removal and before a re-insertion, you will need to update the new HCA port GUID information in the Gateway Switch config file. For information on how to update the configuration, see the *Sun Network QDR Infiniband Gateway Switch Administration Guide*.

1 Become an administrator on the system with multiple IB HCAs.

2 Attempt to unconfigure the attachment point associated with the EoIB datalink.

```
# c f g a d m - c u n c o n f i g u r e P C I - E M 0
c f g a d m : C o m p o n e n t s y s t e m i s b u s y , t r y a g a i n : u n c o n f i g u r e f a i l e d
```

3 Determine the EoIB interface that is preventing the unconfigure operation.

```
# d m e s g | t a i l | g r e p ' f a i l i n g H C A d e t a c h '
A u g 2 3 1 2 : 3 7 : 2 0 e o i b : [ I D 5 3 0 7 9 5 k e r n . w a r n i n g ] W A R N I N G : e o i b 0 s t i l l i n u s e ,
f a i l i n g H C A d e t a c h
```

4 Determine if eoib0 is in use because an IP interface exists on it.

```
# i p a d m s h o w - i f e o i b 0
i p a d m : c a n n o t g e t i n f o r m a t i o n f o r i n t e r f a c e ( s ) : N o s u c h i n t e r f a c e
```

5 If no IP interface over eoib0 exists, check to see if any VNIC is present that is causing eoib0 to be busy.

```
# d l a d m s h o w - v n i c
L I N K O V E R S P E E D M A C A D D R E S S M A C A D D R T Y P E V I D
e v n i c 0 e o i b 0 1 0 0 0 0 2 : 8 : 2 0 : e 5 : 5 6 : 9 9 r a n d o m 0
```

6 Delete the VNIC over eoib0.

```
# d l a d m d e l e t e - v n i c e v n i c 0
```

7 Retry the unconfigure operation.

```
# c f g a d m - c u n c o n f i g u r e P C I - E M 0
```

▼ How to Reconfigure and Restore an EoIB Interface After Hot Removal

1 Become an administrator on the Sun Network QDR Gateway Switch.

2 Edit the BXM config file on the Sun Network QDR Gateway Switch and replace all occurrences of the old-HCA port GUIDs with the port GUIDs of the new HCA replacing it.

- 3 **Restart BXM.**
- 4 **Become an administrator on the Oracle Solaris system.**
- 5 **Hot insert the new HCA on the Solaris system and use the `c f gadm` command to reconfigure the attachment point originally associated with the EoIB datalink.**
- 6 **Confirm the restoration of the original EoIB datalink.**

```
# dladm show-link | grep eoib0
eoib0          phys          1500    unknown    --
```

- 7 **If any VNICs had to be explicitly deleted during a previous unconfigure operation, re-create them now.**

```
# dladm create-vnic -l eoib0 evnic0
```

- 8 **Confirm the presence of the VNIC over eoib0.**

```
# dladm show-vnic
LINK      OVER      SPEED  MACADDRESS      MACADDRTYPE      VID
evnic0    eoib0    10000  2:8:20:e5:56:99  random           0
```

- 9 **If any IP addresses had to be explicitly deleted on eoib0 during a previous unconfigure operation, recreate them by using the `ipadm` command.**

Configuring an IB HCA

Invoke the bus-specific `c f gadm` plugin to configure the HCA. The exact details are beyond the scope of this chapter.

▼ How to Update the IB P_key Tables

If the P_key table information of an HCA ports changes, for example, additional P_keys are enabled or disabled, InfiniBand Transport Framework (IBTF) and IBDM need to be notified so that their internal P_key databases are updated. The `c f gadm` command helps update the P_key databases of IBTF and IBDM. For more information, see [ibtl\(7D\)](#) and [ibdm\(7D\)](#).

- 1 **Become an administrator.**
- 2 **Update the P_key tables.**

For example:

```
# c f gadm -x update_pkey_tbls -y ib
```

▼ How to Display IB Communication Services

Use the following steps to display the communication services that are currently in use by the IBTF.

- 1 **Become an administrator.**
- 2 **Display IB communication services.**

For example:

```
# cfgadm -x list_services ib
Port communication services:
    srp
VPPA communication services:
    ibd
HCA_SVC communication services:
    hnfs
```

▼ How to Add a VPPA Communication Service

Use the following steps to add a new VPPA communication service.

Similar steps can be used to add a new HCA_SVC or a port communication service.

- 1 **Become an administrator.**
- 2 **Add a new VPPA communication service.**
- 3 **Verify that the new service has been added.**

For example:

```
# cfgadm -o comm=vppa,service=new -x add_service ib
```

For example:

```
# cfgadm -x list_services ib
Port communication services:
    srp
VPPA communication services:
    ibd
    new
HCA_SVC communication services:
    nfs_service
```

▼ How to Remove an Existing IB Port, HCA_SVC, or a VPPA Communication Service

Use the following steps to delete an existing IB Port, HCA_SVC, or a VPPA communication service.

1 Become an administrator.

2 Remove a VPPA communication service.

For example:

```
# cfgadm -o comm=vppa,service=new -x delete_service ib
```

3 Verify that the communication service has been removed.

For example:

```
# cfgadm -x list_services ib
Port communication services:
    srp
VPPA communication services:
    ibd
HCA_SVC communication services:
    hnfs
```

Using the uDAPL Application Interface With InfiniBand Devices

User Direct Access Programming Library (uDAPL) is a standard API that promotes data center application data messaging performance, scalability, and reliability over Remote Direct Memory Access (RDMA) capable interconnects such as InfiniBand. The uDAPL interface is defined by the DAT collaborative. For more information about the DAT collaborative, go to the following site:

<http://www.datcollaborative.org>

The Oracle Solaris release provides the following uDAPL features:

- A standard DAT registry library, `libdat`. For more information, see [libdat\(3LIB\)](#).
- A standard service provider registration file, `dat.conf`. For more information, see [dat.conf\(4\)](#).
- Support for multiple service providers so that each provider specifies their own uDAPL library path, version number, and so on, in their own `service_provider.conf` file. For more information, see, [service_provider.conf\(4\)](#).

- An administrative tool, the `datadm` command, to configure `dat.conf`. For more information, see [datadm\(1M\)](#).
- A new resource control property, `project.max-device-locked-memory`, to regulate the amount of locked down physical memory.
- A naming scheme that uses either IPv4 or IPv6 addresses that leverage the IP infrastructure, such as ARP in IPv4 and neighbor discovery in IPv6, for address resolution. The Solaris uDAPL Interface Adapter directly maps to an IPoIB device instance.
- Support for the standard Address Translation Scheme that is used by the DAT collaborative community.
- A uDAPL service provider library to support the hermon Host Channel Adapter with automatic registration to the `dat.conf` registration file. For more information, see [hermon\(7D\)](#).
- Supports both SPARC platform and x86 platforms.

▼ How to Enable uDAPL

- 1 **Become an administrator.**
- 2 **Confirm that the following packages are installed. Or, install them, if needed.**
 - `driver/infiniband/connectx` – InfiniBand Framework
 - `system/io/infiniband/ip-over-ib` – IP over InfiniBand
- 3 **Create the IPoIB interfaces.**

For example:

```
# ipadm create-ip ibd1
# ipadm create-addr -T static -a 192.168.0.1/24 ibd1/ipv4
# datadm -a /usr/share/dat/ABCudapl.t.conf
```

Updating the DAT Static Registry

You can use the `datadm` command to maintain the DAT static registry, the `dat.conf` file. For more information about this file, see [dat.conf\(4\)](#).

The `datadm` command can also be used to register or unregister a service provider to the `dat.conf` file. For more information, see [datadm\(1M\)](#).

When IPoIB interface adapters are added or removed, run the `datadm` command to update the `dat.conf` file to reflect the current state of the system. A new set of interface adapters for all the service providers that are currently installed will be regenerated.

▼ How to Update the DAT Static Registry

- 1 Become an administrator.
- 2 Update the DAT static registry after you add or remove IBoIP interface adapters from the system.

```
# datadm -u
```

- 3 Display the updated DAT static registry.

```
# datadm
```

▼ How to Register a Service Provider in the DAT Static Registry

- 1 Become an administrator.
- 2 Update the DAT static registry after you add the vendor's service provider for the Host Channel Adapter.

```
# datadm -a /usr/share/dat/ABCudaplt.conf
```

- 3 Display the updated DAT static registry.

```
# datadm -v
```

▼ How to Unregister a Service Provider from the DAT Static Registry

- 1 Become an administrator.
- 2 Update the DAT static registry after you remove the vendor's service provider for the Host Channel Adapter from the system.

```
# datadm -r /usr/share/dat/ABCudaplt.conf
```

- 3 Display the updated DAT static registry.

```
# datadm -v
```

Administering IPoB Devices (dLadm)

One physical data link is created by default, per port, per HCA. The physical links can be used as administrative and observability data points. You can create IB partition links over the physical data links, similar to a creating VNICs over a NIC. Keep in mind that the physical data links are not used for data transfers, so plumbing and assigning an IB address is not supported on these links. Data is transferred on the partition data links.

For information about configuring your network components, see *System Administration Guide: IP Services*.

▼ How to Display Physical Data Link Information

Use the `dladm show-phys` command to display physical data link information on your system. The physical link state directly corresponds to the IB HCA port state.

- 1 **Become an administrator.**
- 2 **Display physical data link information.**

For example, display information about `ibp0` — `ibp3` data links.

```
# dladm show-phys
LINK      MEDIA      STATE      SPEED      DUPLEX      DEVICE
ibp0     Infiniband  up         8000       unknown    ibp0
ibp1     Infiniband  down       8000       unknown    ibp1
ibp2     Infiniband  down       8000       unknown    ibp2
ibp3     Infiniband  up         8000       unknown    ibp3
```

For example, use the `show-ib` subcommand to display only the physical links, port GUID, port# HCA GUID, and `P_Key` present on the port at the time the command is running.

```
# dladm show-ib
LINK      HCAGUID      PORTGUID      PORT      STATE      P_Key
ibp0     2C9020040041C 2C9020040041D 1         up         FFFF, 8001
ibp1     2C9020040041C 2C9020040041E 2         down       FFFF
ibp2     3BA0001004E14 3BA0001004E15 1         down       FFFF
ibp3     3BA0001004E14 3BA0001004E16 2         up         FFFF, 8001
```

▼ How to Create IB Partition Links

IB partition data links can be created on top of IB physical links, one per each `P_Key` on the port. The partition data links are used for data transfers.

- 1 **Become an administrator.**
- 2 **Create new IB partition links.**

For example, an IB partition link for `P_Key 0x8001` is created on top of `ibp0` physical data link.

```
# dladm create-part -l ibp0 -P 0x8001 p8001.ibp0
```

The above command succeeds if the port is up, the `P_Key` is present on the port, and IPoIB is successfully initialized.

For example, an IB partition link for the P_Key 0x9000 is created on top of ibp2.

```
# dladm create-part -f -l ibp2 -P 0x9000 p9000.ibp2
```

Note – P_Key 0x9000 is not configured on the port. The force option (-f) option enables you to create an IB partition link, even when the P_Key is not present or the port is down. The link state is marked as down. The link state transitions to up when P_Key is added to the port and port is activated.

3 Display the IB partition link information.

For example:

```
# dladm show-part
LINK      P_Key    OVER      STATE     FLAGS
p8000.ibp0 8001     ibp0      unknown   ----
p9000.ibp2 9000     ibp2      unknown   f---
```

The IB partition link state is up when the link is ready to use, after it is plumbed. The partition link state is down under the following conditions:

- HCA port is down
- P_Key is absent
- Broadcast group is absent

4 Plumb and assign an IP address to an IB partition link.

For example:

```
# ipadm create-ip p9000.ibp2
# ipadm create-addr -T static -a 1.1.1.1 p9000.ibp2/ipv4
# ipadm show-addr
```

▼ How to Display IB Partition Link Information

1 Become an administrator.

2 Display the IB partition link information.

For example:

```
# dladm show-part
LINK      P_Key    OVER      STATE     FLAGS
p8000.ibp0 8001     ibp0      unknown   ----
p9000.ibp2 9000     ibp2      unknown   f---
```

▼ How to Remove an IB Partition Link

1 Become an administrator.

2 Remove an IB partition link.

In this example, partition link p8001.ibp0 is removed.

```
# dladm delete-part p8001.ibp0
```

3 Review the partition link information.

```
# dldam show-part
LINK      P_Key    OVER    STATE    FLAGS
p9000.ibp2  9000    ibp2    down    f---
```

Monitoring and Troubleshooting IB Devices

In the Oracle Solaris 11 release, new commands and utilities enable you to manage the IB fabric more effectively. These commands are included in the `system/io/infiniband/open-fabrics` package and the man pages are installed automatically when the `open-fabrics` package is installed. For example:

```
% man rping
Reformatting page. Please Wait... done

librdmacm                                RPING(1)

NAME
    rping - RDMA CM connection and RDMA ping-pong test.

SYNOPSIS
    rping -s [-v] [-V] [-d] [-P] [-a address] [-p port]
           [-C message_count] [-S message_size]
    rping -c [-v] [-V] [-d] -a address [-p port]
           [-C message_count] [-S message_size]
.
.
.
```

The following new commands and utilities provide the ability to list and query IB devices, diagnose and trouble shoot IB fabric issues, and measure IB performance.

TABLE 6-1 General IB Monitoring Commands

Command	Description
<code>ibv_asyncwatch</code>	Monitors InfiniBand asynchronous events
<code>ibv_devices</code> or <code>ibv_devinfo</code>	Lists InfiniBand devices or device information

TABLE 6-1 General IB Monitoring Commands (Continued)

Command	Description
ibv_rc_pingpong, ibv_srq_pingpong, or ibv_ud_pingpong	Tests node to node connectivity by using RC connection, SRQs, or UD connection
mkey	Tests RDMA CM multicast setup and simple data transfer
rping	Tests RDMA CM connection and attempts RDMA ping-pong
ucmatose	Tests RDMA CM connection and attempts simple ping-pong
udaddy	Tests RDMA CM datagram setup and attempts simple ping-pong

TABLE 6-2 General IB Performance Testing Commands

Command	Description
rdma_bw or rdma_lat	Tests RDMA write transactions for streaming bandwidth or latency
ib_read_bw or ib_read_lat	Tests RDMA read transactions for bandwidth or latency
ib_send_bw or ib_send_lat	Tests RDMA send transactions for bandwidth or latency
ib_write_bw or ib_write_bw_postlist	Tests RDMA write transactions for bandwidth that displays one I/O request at a time or post list bandwidth that displays a list of I/O requests
ib_write_lat	Tests RDMA write transactions for latency
ib_clock_test	Tests accuracy of system clock
qperf	Measures socket and RDMA performance

TABLE 6-3 RDS Monitoring and Testing Tools

Command	Description
rds-info	Displays RDS kernel module information
rds-ping	Determines if remote node over RDS is reachable
rds-stress	Sends message between processes over RDS sockets

TABLE 6-4 Fabric Diagnostic Tools

Command	Description
ibdiagnet	Performs diagnostic check of the entire fabric
ibaddr	Queries InfiniBand address or addresses
ibnetdiscover	Discovers remote InfiniBand topology

TABLE 6-4 Fabric Diagnostic Tools (Continued)

Command	Description
<code>ibping</code>	Validates connectivity between IB nodes
<code>ibportstate</code>	Queries physical port state and link speed of an IB port
<code>ibroute</code>	Displays InfiniBand switch forwarding tables
<code>ibstat</code> or <code>ibsysstat</code>	Query status of InfiniBand device or devices or the status of a system on an IB address
<code>ibtracert</code>	Traces an IB path
<code>perfquery</code> or <code>saquery</code>	Queries IB port counters or sIB subnet administration attributes
<code>sminfo</code>	Queries IB SMInfo attribute
<code>smpquery</code> or <code>smpdump</code>	Queries or dumps IB subnet management attributes
<code>ibcheckerrors</code> or <code>ibcheckerrs</code>	Validates IB port (or node) or IB subnet and reports errors
<code>ibchecknet</code> , <code>ibchecknode</code> , or <code>ibcheckport</code>	Validates IB subnet, node, or port and reports errors
<code>ibcheckportstate</code> , <code>ibcheckportwidth</code> , <code>ibcheckstate</code> , or <code>ibcheckwidth</code>	Validates IB port that are link up but not active, ports for 1x (2.0 Gbps) link width, ports in IB subnet that are link up but not active, or 1x links in IB subnet
<code>ibclearcounters</code> or <code>ibclearerrors</code>	Clears port counters or error counters in IB subnet
<code>ibdatacounters</code> or <code>ibdatacounts</code>	Queries for data counters in IB subnet or IB port data counters
<code>ibdiscover.pl</code>	Annotates and compares IB topology
<code>ibhosts</code>	Displays IB host nodes in topology
<code>iblinkinfo.pl</code> or <code>iblinkinfo</code>	Displays link information for all links in the fabric
<code>ibnodes</code>	Displays IB nodes in topology
<code>ibprintca.pl</code>	Displays either the CA specified or the list of CAs from the <code>ibnetdiscover</code> output
<code>ibprintrt.pl</code>	Displays either only the router specified or a list of routers from the <code>ibnetdiscover</code> output
<code>ibprintswitch.pl</code>	Displays either the switch specified or a list of switches from the <code>ibnetdiscover</code> output
<code>ibqueryerrors.pl</code>	Queries and report non-zero IB port counters
<code>ibrouters</code>	Displays IB router nodes in topology
<code>ibstatus</code>	Queries basic status of IB devices

TABLE 6-4 Fabric Diagnostic Tools (Continued)

Command	Description
<code>ibswitches</code>	Displays IB switch nodes in topology
<code>ibswportwatch.pl</code>	Polls the counters on the specified switch or port and report rate of change information
<code>set_nodedesc.sh</code>	Sets or displays node description string for IB Host Controller Adapters (HCA)s
<code>dump2psl.pl</code>	Dumps PSL file based on opensm output file that is used for credit loop checking
<code>dump2slvl.pl</code>	Dumps SLVL file based on opensm output file that is used for credit loop checking
<code>ibis</code>	An extended TCL shell for IB management inband services

Managing Disks (Overview)

This chapter provides overview information about Oracle Solaris disk slices and introduces the format utility.

This is a list of overview information in this chapter:

- [“What's New in Disk Management?” on page 159](#)
- [“Where to Find Disk Management Tasks” on page 162](#)
- [“Overview of Disk Management” on page 162](#)
- [“Partitioning a Disk” on page 171](#)

For instructions on how to add a disk to your system, see [Chapter 10, “Setting Up Disks \(Tasks\)”](#) or [“x86: Setting Up Disks for ZFS File Systems \(Task Map\)” on page 215](#).

What's New in Disk Management?

The following disk management features are new in this Oracle Solaris release:

- [“Support for Booting From EFI \(GPT\) Labeled Disks” on page 159](#)
- [“Installation Support on Large Disks” on page 160](#)
- [“Advanced Format Disk Support ” on page 161](#)

Support for Booting From EFI (GPT) Labeled Disks

Oracle Solaris installation features can install an EFI (GPT) disk label on a ZFS root pool disk or disks by using DVD, USB, and automated installation methods. UEFI firmware support and the introduction of GRUB 2 provides the ability to boot from a GPT labeled disk. This means that you can use whole disks for the root pool disk or disks on SPARC systems with GPT enabled firmward and on most x86 based systems. Otherwise, systems are installed with a VTOC (SMI) disk label on the root pool disk. For example:

```
# zpool status rpool
pool: rpool
state: ONLINE
scan: none requested
config:

    NAME      STATE    READ WRITE CKSUM
    rpool     ONLINE   0     0     0
    c2t0d0    ONLINE   0     0     0

errors: No known data errors
```

For an example of a EFI (GPT) disk label on a root pool disk, see [Example 9–3](#).

- This disk specification and label support is identified in the Oracle Solaris administration documentation as EFI (GPT).
- SPARC systems require an updated OBP to boot from an EFI (GPT) labeled disk
- An x86 based system that supports GRUB 2 boots from an EFI (GPT) labeled disk.
- The `zpool` command has been enhanced to support EFI (GPT) labels, so that if you need to recreate a root pool or create an alternate root pool after the system is installed, you can do so with the `zpool create -B` command. This new command option creates the required slices and information that is needed for booting.

```
# zpool create -B rpool2 c1t1d0
```

- If you need to replace a disk in a root pool that has an EFI (GPT) labeled disk by using the `zpool replace` command, you also need to reinstall the boot loader. For example:

```
# zpool replace rpool c0t0d0 c1t0d0
# bootadm install-bootloader
```

- The EFI label from previous Oracle Solaris releases is still supported.

Installation Support on Large Disks

New Oracle Solaris installations are no longer limited to the first 2 TiB of the disk on x86 platforms. Oracle Solaris now uses EFI (GPT) partitioning for new installations to enable all of the disk space on the boot device to be used. On x86 platforms, large disk installation is supported through the introduction of GRUB 2 as the default boot loader.

- On SPARC platforms, an OBP update is required. The SPARC boot loader remains unchanged.
- On x86 platforms, large disk installation is supported through the introduction of GRUB 2 as the default system boot loader.

For more information, see [Chapter 2, “Administering the GRand Unified Bootloader \(Tasks\)”](#) in *Booting and Shutting Down Oracle Solaris 11.1 Systems*.

Advanced Format Disk Support

Previous Oracle Solaris releases support disks with a physical block size and a logical block size of 512 bytes. This is the traditional disk block size that is an industry standard.

Currently, disk manufacturers are providing larger capacity disks, also known as advanced format (AF) disks, which is a general term that describes a hard disk drive that exceeds a 512-byte block size.

AF disks are generally in the 4-KB block size range, but vary as follows:

- 4-KB native disk (4kn) – Has a physical and logical block size of 4 KB
- 512-byte emulation (512e) – Has a physical block size of 4 KB but reports a logical block size of 512 bytes

For comparison purposes, Oracle Solaris introduces the 512-byte native (512n) disk term, which is a traditional disk with 512-byte block size.

Oracle Solaris releases support advanced format disks, in addition to traditional 512n disks, in the following ways:

- Oracle Solaris 10 and Oracle Solaris 11 support 4kn and 512e disks for non-root ZFS file systems.
- Oracle Solaris 11.1 provides installation and boot support for 512e devices.

Review the following considerations before purchasing advanced format drives to be used on an Oracle Solaris system:

- Confirm with your device manufacturer that their 512e devices have a power-safe feature to prevent data loss after a power failure when data is still in transit. For more information, see [Oracle Solaris 11.1 Release Notes](#).
- Installation and boot support is not provided on AF disks in Oracle Solaris 10 and Oracle Solaris 11, but they can be used for non-root ZFS file systems.
- Installation and boot support for 4kn devices is not available in Oracle Solaris 11.1.
- Performance is not optimal if 512e and 4kn disks are mixed in existing ZFS storage pools that also contain 512n disks. Performance is best if a new ZFS storage pool is created with all AF disks.

Where to Find Disk Management Tasks

Use these references to find step-by-step instructions for managing disks.

Disk Management Task	For More Information
Format a disk and examine a disk label.	Chapter 9, “Administering Disks (Tasks)”
Add a new disk to a SPARC system.	Chapter 10, “Setting Up Disks (Tasks)”
Add a new disk to an x86 system.	“x86: Setting Up Disks for ZFS File Systems (Task Map)” on page 215
Hot-plug a SCSI or PCI disk.	Chapter 4, “Dynamically Configuring Devices (Tasks)”

Overview of Disk Management

Managing disks in the Oracle Solaris OS usually involves setting up the system and running the Oracle Solaris installation program to create the appropriate disk slices and file systems and to install the Oracle Solaris OS. Occasionally, you might need to use the `format` utility to add a new disk drive or replace a defective disk drive.

The following information is described in this section:

- “Disk Terminology” on page 162
- “About Disk Labels” on page 163
- “EFI (GPT) Disk Label” on page 163
- “About Disk Slices” on page 167
- “`format` Utility” on page 168

Disk Terminology

Before you can effectively use the information that is described in this section, you should be familiar with basic disk architecture. In particular, you should be familiar with the terms in the following table.

Disk Term	Description
Track	A concentric ring on a disk that passes under a single stationary disk head as the disk rotates.
Cylinder	The set of tracks with the same nominal distance from the axis about which the disk rotates.

Disk Term	Description
Sector	Section of each disk platter.
Block	A data storage area on a disk.
Disk controller	A chip and its associated circuitry that controls the disk drive.
Disk label	Part of the disk, usually starting from first sector, that contains disk geometry and partition information.
Device driver	A kernel module that controls a physical (hardware) or virtual device.

For additional information, see the product information from your disk's manufacturer.

About Disk Labels

A special area of every disk is set aside for storing information about the disk's controller, geometry, and slices. This information is called the disk's *label*. Another term that is used to describe the disk label is the *VTOC (Volume Table of Contents)* on a disk with a VTOC label. To *label* a disk means to write slice information onto the disk. You usually label a disk after you change its slices or partitions.

The Oracle Solaris release supports the following two disk labels:

- SMI – The traditional VTOC label for disks that are less than 2 TB in size.
- EFI – Provides support for disks that are larger than 2 TB. The Extensible Firmware Interface GUID Partition Table (EFI GPT) disk label is also available for disks less than 2 TB.

If you fail to label a disk after you create slices, the slices will be unavailable because the OS has no way of “knowing” about the slices.

EFI (GPT) Disk Label

The EFI label provides support for physical disks and virtual disk volumes that are greater than 2 TB in size. This release also includes disk utilities for managing disks greater than 2 TB in size.

Starting in Oracle Solaris 11.1, the system is installed with an EFI (GPT) labeled disk on SPARC systems with GPT enabled firmware and x86 systems by default. For more information, see [“Installing a System With an EFI-Labeled Disk”](#) on page 166.

The following file system products support file systems that are greater than 1 TB in size:

- The Oracle Solaris ZFS file system supports file systems that are greater than 1 TB in size.
- The legacy Solaris Volume Manager software can also be used to manage disks that are greater than 1 TB, but a root disk that is managed by Solaris Volume Manager cannot be used for booting the Oracle Solaris 11 release. For information on using Solaris Volume Manager, see *Solaris Volume Manager Administration Guide*.

You can use the `format -e` command to apply an EFI label to a disk, if the system is running a supported Oracle Solaris release. However, you should review the important information in “[Restrictions of the EFI Disk Label](#)” on page 165 before attempting to apply an EFI label.

You can also use the `format -e` command to reapply a VTOC label, if the EFI label is no longer needed. For example:

```
# format -e
Specify disk (enter its number): 2
selecting c0t5d0
[disk formatted]
.
.
.
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 0
Warning: This disk has an EFI label. Changing to SMI label will erase all
current partitions.
Continue? yes
Auto configuration via format.dat[no]?
Auto configuration via generic SCSI-2[no]?
format> quit
```



Caution – Keep in mind that changing disk labels will destroy any data on the disk.

When using the `format -e` command on an EFI (GPT) labeled disk, the partition menu displays 128 partitions (slices), but only 7 partitions are usable.

Comparison of the EFI Label and the VTOC Label

The EFI disk label differs from the VTOC disk label in the following ways:

- Provides support for disks greater than 2 terabytes in size.
- Provides usable slices 0-6, where partition 2 is just another slice.
- Partitions (or slices) cannot overlap with the primary or backup label, nor with any other partitions. The size of the EFI label is usually 34 sectors, so partitions usually start at sector 34. This feature means that no partition can start at sector zero (0).

- EFI (GPT) labels do not use the notion of geometry. The partitions are defined based on logical blocks.
- Some information that was stored in the alternate cylinder area is now stored in the last two cylinders of a disk or Solaris partition.
- If you use the `format` utility to change partition sizes, the unassigned partition tag is assigned to partitions with sizes equal to zero. By default, the `format` utility assigns the `usr` partition tag to any partition with a size greater than zero. You can use the partition change menu to reassign partition tags, after the partitions are changed. However, you cannot change a partition with a non-zero size to the unassigned partition tag.

Restrictions of the EFI Disk Label

Keep the following restrictions in mind when determining whether using disks greater than 2 terabytes is appropriate for your environment:

- Layered software products that are intended for systems with VTOC-labeled disks might be incapable of accessing a disk with an EFI disk label.
- On x86 based systems, you can use the `fdisk` command on a disk with an EFI label that is greater than 2 TB in size.
- Use the `format` utility to partition disks with EFI labels.
- The EFI specification prohibits overlapping partitions. The entire disk is represented by `cxydz`.
- The EFI disk label provides information about disk or partition sizes in sectors and blocks, but not in cylinders and heads.
- The following `format` options are either not supported or are not applicable to disks with EFI labels:
 - The `save` option is not supported, because disks with EFI labels do not need an entry in the `format.dat` file.
 - The `backup` option is not applicable.

x86: Support for EFI-Labeled Disks

Oracle Solaris support for the EFI disk label is available on x86 systems. Use the following command to add an EFI label on an x86 system:

```
# format -e
> [0] SMI Label
> [1] EFI Label
> Specify Label type[0]: 1
> WARNING: converting this device to EFI labels will erase all current
> fdisk partition information. Continue? yes
```

Previous label information is not converted to the EFI disk label.

You will have to recreate the label's partition information manually with the `format` command. You cannot use the `fdisk` command on a disk with an EFI label that is 2 terabytes in size. If the `fdisk` command is run on a disk that is greater than 2 TB in size to create a Solaris partition, the Solaris partition is limited to 2 TB. For more information about EFI disk labels, see the preceding section.

Installing a System With an EFI-Labeled Disk

In Oracle Solaris 11, a root pool disk must have an SMI label. The installation utilities automatically relabel any disk that is selected as a root pool disk with an SMI label.

In Oracle Solaris 11.1, in most cases, when the system is installed, an EFI (GPT) label is applied automatically to the root pool disk on SPARC systems with GPT enabled firmware and x86 based systems. For example:

```
# zpool status rpool
pool: rpool
state: ONLINE
scan: none requested
config:

    NAME      STATE    READ WRITE CKSUM
    rpool     ONLINE   0     0     0
    c8t2d0    ONLINE   0     0     0
```

After installation on an x86 based system, a root pool disk might look similar to the following:

```
# prtvtoc /dev/dsk/c8t2d0
* /dev/dsk/c8t2d0 partition map
*
* Dimensions:
*   512 bytes/sector
* 143374738 sectors
* 143374671 accessible sectors
*
* Flags:
*   1: unmountable
*  10: read-only
*
* Unallocated space:
*   First Sector      Last
*   Sector  Count     Sector
*     34      222      255
*
* Partition Tag  Flags  First Sector      Sector  Count     Last Sector  Mount Directory
*     0     24   00      256     524288     524543
*     1     4    00    524544  142833777  143358320
*     8    11   00  143358321  16384  143374704
root@sys-04:~#
```

On an x86 system, in addition to the traditional partition 8, a small partition 0 is created to contain the boot loader. Similar to partition 8, this slice requires no administration and should be left alone. The root file system is contained in partition 1.

The `zpool` command has been modified to create a new root pool disk label automatically, if you need to recreate a root pool after the system is installed. For more information, see [Chapter 4, “Managing ZFS Root Pool Components,”](#) in *Oracle Solaris 11.1 Administration: ZFS File Systems*.

Managing Disks With EFI Labels

Use the following table to locate information on managing disks with EFI labels.

Task	For More Information
If the system is not installed, install it.	<i>Installing Oracle Solaris 11.1 Systems</i>
The system is already installed, but the root pool disk was damaged or needs to be replaced.	“SPARC: How to Set Up a Disk for a ZFS Root File System” on page 209 or “x86: How to Set Up a Disk for a ZFS Root File System” on page 217
The system is already installed, but you need to set up a disk for a non-root pool.	“SPARC: How to Set Up a Disk for a ZFS Non-Root File System” on page 214 or “x86: How to Set Up a Disk for a ZFS Non-Root File System” on page 228

Troubleshooting Problems With EFI Disk Labels

Use the following error messages and solutions to troubleshoot problems with EFI-labeled disks.

Cause

Boot a system running a SPARC or x86 kernel with a disk greater than 1 terabyte.

Error Message

```
Dec  3 09:12:17 holoship scsi: WARNING: /sbus@a,0/SUNW,socal@d,10000/
sf@1,0/ssd@w50020f23000002a4,0 (ssd1):
Dec  3 09:12:17 holoship corrupt label - wrong magic number
```

Cause

You attempted to add a disk to a system running an older Solaris release.

Solution

Add the disk to a system running the Solaris release that supports the EFI disk label.

About Disk Slices

Files that are stored on a disk are contained in file systems. Each file system on a disk is assigned to a *slice*, which is a group of sectors that are set aside for use by that file system. Each disk slice appears to the Oracle Solaris OS (and to the system administrator) as though it were a separate disk drive.

For information about file systems, see [Chapter 14, “Managing File Systems \(Overview\).”](#)

Note – Slices are sometimes referred to as *partitions*. Certain interfaces, such as the `format` utility, refer to slices as partitions.

When setting up slices, remember these rules:

- Each disk slice holds only one file system.
- No file system can span multiple slices.

Using Raw Data Slices

The disk label is stored in block 0 of each disk. So, third-party database applications that create raw data slices must not start at block 0. Otherwise, the disk label will be overwritten, and the data on the disk will be inaccessible.

Do not use the following areas of the disk for raw data slices, which are sometimes created by third-party database applications:

- Block 0 where the disk label is stored
- Slice 2, which represents the entire disk with a VTOC label

format Utility

Read the following overview of the `format` utility and its uses before proceeding to the “how-to” or reference sections.

The `format` utility is a system administration tool that is used to prepare hard disk drives for use on your Oracle Solaris system.

The following table describes the features and associated benefits of the `format` utility.

TABLE 7-1 Features and Benefits of the `format` Utility

Feature	Benefit
Searches your system for all attached disk drives	Reports on the following: <ul style="list-style-type: none">▪ Target location▪ Disk geometry▪ Whether the disk is formatted▪ If the disk has mounted partitions
Retrieves disk labels	Convenient for repair operations
Repairs defective sectors	Allows administrators to repair disk drives with recoverable errors instead of sending the drive back to the manufacturer
Formats and analyzes a disk	Creates sectors on the disk and verifies each sector

TABLE 7-1 Features and Benefits of the format Utility (Continued)

Feature	Benefit
Partitions a disk	Divides a disk into slices or partitions. ZFS file systems do not correspond to disk slices or partitions, except for the ZFS root pool.
Labels a disk	Writes disk name and configuration information to the disk for future retrieval (usually for repair operations)

The format utility options are described in [Chapter 13, “The format Utility \(Reference\).”](#)

When to Use the format Utility

Disk drives are partitioned and labeled by the Oracle Solaris installation utility when you install Oracle Solaris. You can use the format utility to do the following:

- Display slice or partition information
- Partition a disk
- Add a disk drive to an existing system
- Format a disk drive
- Label a disk
- Repair a disk drive
- Analyze a disk for errors

The main reason a system administrator uses the format utility is to partition a disk. These steps are covered in [Chapter 10, “Setting Up Disks \(Tasks\),”](#) and [“x86: Setting Up Disks for ZFS File Systems \(Task Map\)”](#) on page 215.

See the following section for guidelines on using the format utility.

Guidelines for Using the format Utility

TABLE 7-2 format Utility Guidelines

Task	Guidelines	For More Information
Format a disk.	<ul style="list-style-type: none"> ■ Any existing data is destroyed when you reformat a disk. ■ The need for formatting a disk drive has decreased as more and more manufacturers ship their disk drives formatted and partitioned. You might not need to use the format utility when you add or replace a disk drive to an existing system. ■ If a disk has been relocated and is displaying many disk errors, you can attempt to relabel it. 	“How to Format a Disk” on page 193 or “How to Label a Disk” on page 197

TABLE 7-2 format Utility Guidelines (Continued)

Task	Guidelines	For More Information
Set up a disk that contains a ZFS root file system.	In a non-redundant configuration, a ZFS root file system data from the damaged disk must be restored from a backup medium. Otherwise, the system will have to be reinstalled by using the installation utility.	“SPARC: How to Set Up a Disk for a ZFS Root File System” on page 209 or “x86: How to Set Up a Disk for a ZFS Root File System” on page 217 , or, if the system must be reinstalled, Installing Oracle Solaris 11.1 Systems
Create a VTOC labeled disk slice for a root pool on a SPARC based system. Or, create an EFI labeled disk partition for a root pool on x86 based system.	<ul style="list-style-type: none"> ▪ The best way to use a ZFS storage pool is by creating a pool with whole disks. ▪ If a disk is intended to be used for a root pool on a SPARC based system, you must create a disk slice. This is long-standing boot limitation. 	“SPARC: How to Create a Disk Slice for a ZFS Root File System” on page 210 or “x86: How to Set Up a Disk for a ZFS Root File System” on page 217
Set up a disk that contains a ZFS non-root file system.	A disk that is used for a non-root ZFS file system usually contains space for user or data files. You can attach or add another disk to a root pool or a non-root pool for more disk space.	“SPARC: How to Set Up a Disk for a ZFS Non-Root File System” on page 214 or “x86: How to Set Up a Disk for a ZFS Non-Root File System” on page 228

Formatting a Disk

In most cases, disks are formatted by the manufacturer or reseller. So, they do not need to be reformatted when you install the drive. To determine if a disk is formatted, use the `format` utility. For more information, see [“How to Determine if a Disk Is Formatted” on page 192](#).

If you determine that a disk is not formatted, use the `format` utility to format the disk.

When you format a disk, you accomplish two steps:

- The disk media is prepared for use.
- A list of disk defects based on a surface analysis is compiled.



Caution – Formatting a disk is a destructive process because it overwrites data on the disk. For this reason, disks are usually formatted only by the manufacturer or reseller. If you think disk defects are the cause of recurring problems, you can use the `format` utility to do a surface analysis. However, be careful to use only the commands that do not destroy data. For details, see [“How to Format a Disk” on page 193](#).

A small percentage of total disk space that is available for data is used to store defect and formatting information. This percentage varies according to disk geometry, and decreases as the disk ages and develops more defects.

Formatting a disk might take anywhere from a few minutes to several hours, depending on the type and size of the disk.

Partitioning a Disk

The following information is described in this section:

- “Partition Table Terminology” on page 171
- “Displaying Partition Table Information” on page 172
- “Using the Free Hog Slice” on page 174

The `format` utility is most often used by system administrators to partition a disk. The steps are as follows:

- Determine which slices are needed.
- Determine the size of each slice or partition.
- Use the `format` utility to partition the disk.
- Label the disk with new partition information.
- Create the file system for each partition.

The easiest way to partition a disk is to use the `modify` command from the partition menu of the `format` utility. The `modify` command enables you to create partitions by specifying the size of each partition, without having to keep track of the starting cylinder boundaries. The `modify` command also keeps tracks of any disk space that remains in the “free hog” slice.

Partition Table Terminology

An important part of the disk label is the *partition table*. The partition table identifies a disk's slices, the slice boundaries (in cylinders), and the total size of the slices. You can display a disk's partition table by using the `format` utility. The following table describes partition table terminology.

TABLE 7-3 Partition Table Terminology

Partition Term	Value	Description
Number	0–7	VTOC – Partitions or slices, numbered 0–7. EFI – Partitions numbered 0–6.
Tag	0=UNASSIGNED 1=BOOT 2=ROOT 3=SWAP 4=USR 5=BACKUP 7=VAR 8=HOME 11=RESERVED	A numeric value that usually describes the file system mounted on this partition.

TABLE 7-3 Partition Table Terminology (Continued)

Partition Term	Value	Description
Flags	wm	The partition is writable and mountable.
	wu rm	The partition is writable and unmountable. This state is the default for partitions that are dedicated for swap areas. (However, the mount command does not check the “not mountable” flag.)
	rm	The partition is read only and mountable.

Partition flags and tags are assigned by convention and require no maintenance.

For more information about displaying the partition table, see the following references:

- “Displaying Partition Table Information” on page 172
- “How to Display Disk Slice Information” on page 195
- “How to Examine a Disk Label” on page 202

Displaying Partition Table Information

The following format utility output shows an example of a partition table from a 74-GB disk with a VTOC label displayed:

Total disk cylinders available: 38756 + 2 (reserved cylinders)

Part	Tag	Flag	Cylinders	Size	Blocks
0	root	wm	3 - 2083	4.00GB	(2081/0/0) 8390592
1	swap	wu	2084 - 3124	2.00GB	(1041/0/0) 4197312
2	backup	wm	0 - 38755	74.51GB	(38756/0/0) 156264192
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	unassigned	wm	0	0	(0/0/0) 0
7	home	wm	3125 - 38755	68.50GB	(35631/0/0) 143664192
8	boot	wu	0 - 0	1.97MB	(1/0/0) 4032
9	alternates	wu	1 - 2	3.94MB	(2/0/0) 8064

partition>

The partition table that is displayed by the format utility contains the following information.

Column Name	Description
Part	Partition or slice number. See Table 7-3 for a description of this column.
Tag	Partition tag. See Table 7-3 for a description of this column.
Flag	Partition flag. See Table 7-3 for a description of this column.

Column Name	Description
Cylinders	The starting and ending cylinder number for the slice. Not displayed on EFI-labeled disks.
Size	The slice or partition size in MB.
Blocks	The total number of cylinders and the total number of sectors per slice. Not displayed on EFI-labeled disks.
First Sector	EFI – The starting block number. Not displayed on VTOC-labeled disks.
Last Sector	EFI – The ending block number. Not displayed on VTOC-labeled disks.

The following is an example of an EFI disk label displayed by using the `prtvtoc` command.

```
# prtvtoc /dev/rdisk/c4t1d0s0
* /dev/rdisk/c4t1d0s0 partition map
*
* Dimensions:
*   512 bytes/sector
* 2576941056 sectors
* 2576940989 accessible sectors
*
* Flags:
*  1: unmountable
* 10: read-only
*
*
* Partition  Tag  Flags      First      Sector      Last
* Partition  Tag  Flags      Sector     Count       Sector  Mount Directory
*   0         2    00         34        629145600   629145633
*   1         4    00    629145634   629145600   1258291233
*   6         4    00  1258291234  1318633404  2576924637
*   8        11    00  2576924638         16384  2576941021
```

The output of the `prtvtoc` command provides information in the following three sections:

- Dimensions
- Flags
- Partition Table

prtvtoc Column Name	Description
Partition	Partition or slice number. For a description of this column, see Table 7-3 .
Tag	Partition tag. For a description of this column, see Table 7-3 .
Flags	Partition flag. For a description of this column, see Table 7-3 .
First Sector	The first sector of the slice or partition.
Sector Count	The total number of sectors in the slice or partition.
Last Sector	The last sector of the slice or partition.

<code>prvtoc</code> Column Name	Description
Mount Directory	The last mount point directory for the file system.

Using the Free Hog Slice

When you use the `format` utility to change the size of one or more disk slices, you designate a temporary slice that will expand and shrink to accommodate the resizing operations.

This temporary slice donates, or “frees,” space when you expand a slice, and receives, or “hogs,” the discarded space when you shrink a slice. For this reason, the donor slice is sometimes called the *free hog*.

The free hog slice exists only during installation or when you run the `format` utility. There is no permanent free hog slice during day-to-day operations.

For information on using the free hog slice, see [“SPARC: How to Create a Disk Slice for a ZFS Root File System” on page 210](#) or [“x86: How to Replace a ZFS Root Pool Disk \(EFI \(GPT\)\)” on page 222](#).

Managing Disk Use (Tasks)

This chapter describes how to optimize disk space by locating unused files and large directories.

This is a list of the information in this chapter:

- “Managing Disk Use (Task Map)” on page 175
- “Displaying Information About Files and Disk Space” on page 176
- “Checking the Size of Files” on page 178
- “Checking the Size of Directories” on page 182
- “Finding and Removing Old or Inactive Files” on page 184

Managing Disk Use (Task Map)

Task	Description	For Instructions
Display information about files and disk space.	Display information about how disk space is used by using the <code>df</code> command.	“How to Display Information About Files and Disk Space” on page 177
Display the size of files.	Display information about the size of files by using the <code>ls</code> command with the <code>-lh</code> options.	“How to Display the Size of Files” on page 179
Find large files.	The <code>ls -s</code> command allows you to sort files by size, in descending order.	“How to Find Large Files” on page 180
Find files that exceed a specified size limit.	Locate and display the names of files that exceed a specified size by using the <code>find</code> command with the <code>-size</code> option and the value of the specified size limit.	“How to Find Files That Exceed a Specified Size Limit” on page 181

Task	Description	For Instructions
Display the size of directories, subdirectories, and files.	Display the size of one or more directories, subdirectories, and files by using the <code>du</code> command.	“How to Display the Size of Directories, Subdirectories, and Files” on page 182
List the newest files.	Display the most recently created or changed files first, by using the <code>ls -t</code> command.	“How to List the Newest Files” on page 184
Find and remove old or inactive files.	Use the <code>find</code> command with the <code>-atime</code> and <code>-mtime</code> options to locate files that have not been accessed for a specified number of days. You can remove these files by using the <code>cat filename</code> command.	“How to Find and Remove Old or Inactive Files” on page 185
Clear out temporary directories.	Locate temp directories, then use the <code>rm -r *</code> command to remove the entire directory.	“How to Clear Out Temporary Directories” on page 186
Find and delete core files.	Find and delete core files by using the <code>find . -name core -exec rm {} \;</code> command.	“How to Find and Delete core Files” on page 187
Delete crash dump files.	Delete crash dump files that are located in the <code>/var/crash/</code> directory by using the <code>rm *</code> command.	“How to Delete Crash Dump Files” on page 187

Displaying Information About Files and Disk Space

This table summarizes the commands available for displaying information about file size and disk space.

Command	Description	Man Page
<code>df</code>	Reports the number of free disk blocks and files	df(1M)
<code>du</code>	Summarizes disk space allocated to each subdirectory	du(1)
<code>find -size</code>	Searches recursively through a directory based on the size specified with the <code>-size</code> option	find(1)
<code>ls -lh</code>	Lists the size of a file in the power of 1024 scaling	ls(1)

▼ How to Display Information About Files and Disk Space

- Display information about how disk space is used by using the `df` command.

```
$ df [directory] [-h] [-t]
```

`df` With no options, lists all mounted file systems and their device names, the number of 512-byte blocks used, and the number of files.

directory Specifies the directory whose file system you want to check.

`-h` Displays disk space in the power of 1024 scaling.

`-t` Displays the total blocks as well as the blocks used for all mounted file systems.

Example 8–1 Displaying Information About File Size and Disk Space

In the following example, all the file systems listed are locally mounted except for `/usr/dist`.

```
$ df
/                (/dev/dsk/c0t0d0s0 ): 101294 blocks  105480 files
/devices         (/devices             ):      0 blocks      0 files
/system/contract (ctfs                 ):      0 blocks 2147483578 files
/proc            (proc                 ):      0 blocks   1871 files
/etc/mnttab      (mnttab               ):      0 blocks      0 files
/etc/svc/volatile (swap                 ): 992704 blocks  16964 files
/system/object   (objfs                ):      0 blocks 2147483530 files
/usr             (/dev/dsk/c0t0d0s6 ): 503774 blocks 299189 files
/dev/fd          (fd                   ):      0 blocks      0 files
/var/run         (swap                 ): 992704 blocks  16964 files
/tmp             (swap                 ): 992704 blocks  16964 files
/opt             (/dev/dsk/c0t0d0s5 ): 23914 blocks  6947 files
/export/home     (/dev/dsk/c0t0d0s7 ): 16810 blocks  7160 files
```

Example 8–2 Displaying File Size Information in 1024 Bytes on a System With a ZFS Root File System

In the following example, file system information for a system is displayed in 1024 bytes.

Filesystem	size	used	avail	capacity	Mounted on
rpool/ROOT/s1008be	67G	4.6G	58G	8%	/
/devices	0K	0K	0K	0%	/devices
ctfs	0K	0K	0K	0%	/system/contract
proc	0K	0K	0K	0%	/proc
mnttab	0K	0K	0K	0%	/etc/mnttab
swap	1.9G	1.5M	1.9G	1%	/etc/svc/volatile
objfs	0K	0K	0K	0%	/system/object
sharefs	0K	0K	0K	0%	/etc/dfs/sharetab
/platform/sun4u-us3/lib/libc_psr/libc_psr_hwcap1.so.1	63G	4.6G	58G	8%	/platform/sun4u-us3/lib/libc_psr.so.1

```

/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr_hwcap1.so.1
63G 4.6G 58G 8% /platform/sun4u-us3/lib/sparcv9/libc_psr.so.1
fd 0K 0K 0K 0% /dev/fd
rpool/ROOT/s1008be/var
67G 73M 58G 1% /var
swap 1.9G 32K 1.9G 1% /tmp
swap 1.9G 40K 1.9G 1% /var/run
rpool/export 67G 20K 58G 1% /export
rpool/export/home 67G 18K 58G 1% /export/home

```

Example 8-3 Displaying Total Number of Blocks and Files Allocated for a File System

The following example shows a list of all mounted file systems, device names, total 512-byte blocks used, and the number of files. The second line of each two-line entry displays the total number of blocks and files that are allocated for the file system.

```

$ df -t
/ (/dev/dsk/c0t0d0s0 ): 101294 blocks 105480 files
total: 509932 blocks 129024 files
/devices (/devices ): 0 blocks 0 files
total: 0 blocks 113 files
/system/contract (ctfs ): 0 blocks 2147483578 files
total: 0 blocks 69 files
/proc (proc ): 0 blocks 1871 files
total: 0 blocks 1916 files
/etc/mnttab (mnttab ): 0 blocks 0 files
total: 0 blocks 1 files
/etc/svc/volatile (swap ): 992608 blocks 16964 files
total: 993360 blocks 17025 files
/system/object (objfs ): 0 blocks 2147483530 files
total: 0 blocks 117 files
/usr (/dev/dsk/c0t0d0s6 ): 503774 blocks 299189 files
total: 6650604 blocks 420480 files
/dev/fd (fd ): 0 blocks 0 files
total: 0 blocks 31 files
/var/run (swap ): 992608 blocks 16964 files
total: 992688 blocks 17025 files
/tmp (swap ): 992608 blocks 16964 files
total: 992688 blocks 17025 files
/opt (/dev/dsk/c0t0d0s5 ): 23914 blocks 6947 files
total: 27404 blocks 7168 files
/export/home (/dev/dsk/c0t0d0s7 ): 16810 blocks 7160 files
total: 18900 blocks 7168 files

```

Checking the Size of Files

You can check the size of files and sort them by using the `ls` command. You can find files that exceed a size limit by using the `find` command. For more information, see the [ls\(1\)](#) and [find\(1\)](#) man pages.

Note – If you run out of space in the `/var` directory, do not symbolically link the `/var` directory to a directory on a file system with more disk space. Doing so, even as a temporary measure, might cause problems for certain daemon processes and utilities.

▼ How to Display the Size of Files

- 1 Change to the directory where the files you want to check are located.
- 2 Display the size of the files.

```
$ ls [-lh] [-s]
```

- l Displays a list of files and directories in long format, showing the sizes in bytes. (See the example that follows.)
- h Scales file sizes and directory sizes into Kbytes, Mbytes, Gbytes, or Tbytes when the file or directory size is larger than 1024 bytes. This option also modifies the output displayed by the `-o`, `-n`, `-@`, and `-g` options to display file or directory sizes in the new format. For more information, see the `ls(1)` man page.
- s Displays a list of the files and directories, showing the sizes in blocks.

Example 8–4 Displaying the Size of Files

The following example shows that the `lastlog` and `messages` files are larger than the other files in the `/var/adm` directory.

```
$ cd /var/adm
$ ls -lh
total 148
drwxrwxr-x  5 adm      adm           512 Nov 26 09:39 acct/
-rw-----  1 uucp     bin           0 Nov 26 09:25 aculog
drwxr-xr-x  2 adm      adm           512 Nov 26 09:25 exacct/
-r--r--r--  1 root     other        342K Nov 26 13:56 lastlog
drwxr-xr-x  2 adm      adm           512 Nov 26 09:25 log/
-rw-r--r--  1 root     root         20K Nov 26 13:55 messages
drwxr-xr-x  2 adm      adm           512 Nov 26 09:25 passwd/
drwxrwxr-x  2 adm      sys           512 Nov 26 09:39 sa/
drwxr-xr-x  2 root     sys           512 Nov 26 09:49 sm.bin/
-rw-rw-rw-  1 root     bin           0 Nov 26 09:25 spellhist
drwxr-xr-x  2 root     sys           512 Nov 26 09:25 streams/
-rw-r--r--  1 root     bin          3.3K Nov 26 13:56 utmpx
-rw-r--r--  1 root     root           0 Nov 26 10:17 vold.log
-rw-r--r--  1 adm      adm           19K Nov 26 13:56 wtmpx
```

The following example shows that the `lpsched.1` file uses two blocks.

```
$ cd /var/lp/logs
$ ls -s
total 2          0 lpsched          2 lpsched.1
```

▼ How to Find Large Files

- 1 Change to the directory that you want to search.
- 2 Display the size of files in blocks from largest to smallest.
 - If the characters or columns for the files are *different*, use the following command to sort a list of files by block size, from largest to smallest.

```
$ ls -l | sort +4rn | more
```

Note that this command sorts files in a list by the character that is in the fourth field, starting from the left.

- If the characters or columns for the files are the *same*, use the following command to sort a list of files by block size, from largest to smallest.

```
$ ls -s | sort -nr | more
```

Note that this command sorts files in a list, starting with the left most character.

Example 8-5 Finding Large Files (Sorting by the Fifth Field's Character)

```
$ cd /var/adm
$ ls -l | sort +4rn | more
-r--r--r-- 1 root  root  4568368 Oct 17 08:36 lastlog
-rw-r--r-- 1 adm   adm   697040 Oct 17 12:30 pacct.9
-rw-r--r-- 1 adm   adm   280520 Oct 17 13:05 pacct.2
-rw-r--r-- 1 adm   adm   277360 Oct 17 12:55 pacct.4
-rw-r--r-- 1 adm   adm   264080 Oct 17 12:45 pacct.6
-rw-r--r-- 1 adm   adm   255840 Oct 17 12:40 pacct.7
-rw-r--r-- 1 adm   adm   254120 Oct 17 13:10 pacct.1
-rw-r--r-- 1 adm   adm   250360 Oct 17 12:25 pacct.10
-rw-r--r-- 1 adm   adm   248880 Oct 17 13:00 pacct.3
-rw-r--r-- 1 adm   adm   247200 Oct 17 12:35 pacct.8
-rw-r--r-- 1 adm   adm   246720 Oct 17 13:15 pacct.0
-rw-r--r-- 1 adm   adm   245920 Oct 17 12:50 pacct.5
-rw-r--r-- 1 root  root  190229 Oct  5 03:02 messages.1
-rw-r--r-- 1 adm   adm   156800 Oct 17 13:17 pacct
-rw-r--r-- 1 adm   adm   129084 Oct 17 08:36 wtmpx
```

Example 8-6 Finding Large Files (Sorting by the Left Most Character)

In the following example, the `lastlog` and `messages.1` files are the largest files in the `/var/adm` directory.

```

$ cd /var/adm
$ ls -s | sort -nr | more
 48 lastlog
 30 messages
 24 wtmpx
 18 pacct
  8 utmpx
  2 vold.log
  2 sulog
  2 sm.bin/
  2 sa/
  2 passwd/
  2 pacct1
  2 log/
  2 acct/
  0 spellhist
  0 aculog
total 144

```

▼ How to Find Files That Exceed a Specified Size Limit

- To locate and display the names of files that exceed a specified size, use the `find` command.

```
$ find directory -size +nnn
```

directory Identifies the directory that you want to search.

`-size +nnn` Is a number of 512-byte blocks. Files that exceed this size are listed.

Example 8-7 Finding Files That Exceed a Specified Size Limit

The following example shows how to find files larger than 400 blocks in the current working directory. The `-print` option displays the output of the `find` command.

```

$ find . -size +400 -print
./Howto/howto.doc
./Howto/howto.doc.backup
./Howto/howtotest.doc
./Routine/routineBackupconcepts.doc
./Routine/routineIntro.doc
./Routine/routineTroublefsck.doc
./.record
./Mail/pagination
./Config/configPrintadmin.doc
./Config/configPrintsetup.doc
./Config/configMailappx.doc
./Config/configMailconcepts.doc
./snapshot.rs

```

Checking the Size of Directories

You can display the size of directories by using the `du` command and options. Additionally, you can find the amount of disk space used by user accounts on local UFS file systems by using the `quot` command. For more information about these commands, see the [du\(1\)](#) and [quot\(1M\)](#) man pages.

▼ How to Display the Size of Directories, Subdirectories, and Files

- Display the size of one or more directories, subdirectories, and files by using the `du` command. Sizes are displayed in 512-byte blocks.

```
$ du [-as] [directory ...]
```

<code>du</code>	Displays the size of each directory that you specify, including each subdirectory beneath it.
<code>-a</code>	Displays the size of each file and subdirectory, and the total number of blocks that are contained in the specified directory.
<code>-s</code>	Displays the total number of blocks that are contained in the specified directory.
<code>-h</code>	Displays the size of each directory in 1024-byte blocks.
<code>-H</code>	Displays the size of each directory in 1000-byte blocks.
<code>[directory ...]</code>	Identifies one or more directories that you want to check. Separate multiple directories in the command-line syntax with spaces.

Example 8–8 Displaying the Size of Directories, Subdirectories, and Files

The following example shows the sizes of two directories:

```
$ du -s /var/adm /var/spool/lp
130    /var/adm
40     /var/spool/lp
```

The following example shows the sizes of two directories and includes the sizes of all the subdirectories and files that are contained within each directory. The total number of blocks that are contained in each directory is also displayed:

```
$ du /var/adm /var/spool/lp
2      /var/adm/exacct
2      /var/adm/log
2      /var/adm/streams
```

```

2      /var/adm/acct/fiscal
2      /var/adm/acct/nite
2      /var/adm/acct/sum
8      /var/adm/acct
2      /var/adm/sa
2      /var/adm/sm.bin
258    /var/adm
4      /var/spool/lp/admins
2      /var/spool/lp/requests/printing....
4      /var/spool/lp/requests
4      /var/spool/lp/system
2      /var/spool/lp/fifos
24     /var/spool/lp

```

The following example shows directory sizes in 1024-byte blocks:

```

$ du -h /usr/share/audio
796K  /usr/share/audio/samples/au
797K  /usr/share/audio/samples
798K  /usr/share/audio

```

▼ How to Display the User Ownership of Local UFS File Systems

1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Display users, directories, or file systems, and the number of 1024-byte blocks used.

```
# quot [-a] [filesystem ...]
```

-a Lists all users of each mounted UFS file system and the number of 1024-byte blocks used.

filesystem Identifies a UFS file system. Users and the number of blocks used are displayed for that file system.

Note – The quot command works only on local UFS file systems.

Example 8–9 Displaying the User Ownership of Local UFS File Systems

In the following example, users of the root (/) file system are displayed. In the subsequent example, users of all mounted UFS file systems are displayed.

```

# quot /
/dev/rdisk/c0t0d0s0:

```

```
43340 root
3142 rimmer
47 uucp
35 lp
30 adm
4 bin
4 daemon

# quot -a
/dev/rdisk/c0t0d0s0 (/):
43340 root
3150 rimmer
47 uucp
35 lp
30 adm
4 bin
4 daemon
/dev/rdisk/c0t0d0s6 (/usr):
460651 root
206632 bin
791 uucp
46 lp
4 daemon
1 adm
/dev/rdisk/c0t0d0s7 (/export/home):
9 root
```

Finding and Removing Old or Inactive Files

Part of the job of cleaning up heavily loaded file systems involves locating and removing files that have not been used recently. You can locate unused files by using the `ls` or `find` commands. For more information, see the [ls\(1\)](#) and [find\(1\)](#) man pages.

Other ways to conserve disk space include emptying temporary directories such as the directories located in `/var/tmp` or `/var/spool`, and deleting core and crash dump files. For more information about crash dump files, refer to [Chapter 1, “Managing System Crash Information \(Tasks\)”](#), in *Troubleshooting Typical Issues in Oracle Solaris 11.1*.

▼ How to List the Newest Files

- List files, displaying the most recently created or changed files first, by using the `ls -t` command.

```
$ ls -t [directory]
```

`-t` Sorts files by latest time stamp first.

`directory` Identifies the directory that you want to search.

Example 8–10 Listing the Newest Files

The following example shows how to use the `ls -tl` command to locate the most recently created or changed files within the `/var/adm` directory. The `su` log file was created or edited most recently.

```
$ ls -tl /var/adm
total 134
-rw----- 1 root   root       315 Sep 24 14:00 su_log
-r--r--r-- 1 root   other     350700 Sep 22 11:04 lastlog
-rw-r--r-- 1 root   bin       4464 Sep 22 11:04 utmpx
-rw-r--r-- 1 adm    adm       20088 Sep 22 11:04 wtmpx
-rw-r--r-- 1 root   other     0 Sep 19 03:10 messages
-rw-r--r-- 1 root   other     0 Sep 12 03:10 messages.0
-rw-r--r-- 1 root   root     11510 Sep 10 16:13 messages.1
-rw-r--r-- 1 root   root     0 Sep 10 16:12 vold.log
drwxr-xr-x 2 root   sys      512 Sep 10 15:33 sm.bin
drwxrwxr-x 5 adm    adm      512 Sep 10 15:19 acct
drwxrwxr-x 2 adm    sys      512 Sep 10 15:19 sa
-rw----- 1 uucp   bin       0 Sep 10 15:17 aculog
-rw-rw-rw- 1 root   bin       0 Sep 10 15:17 spellhist
drwxr-xr-x 2 adm    adm      512 Sep 10 15:17 log
drwxr-xr-x 2 adm    adm      512 Sep 10 15:17 passwd
```

▼ How to Find and Remove Old or Inactive Files

1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Find files that have not been accessed for a specified number of days and list them in a file.

```
# find directory -type f[-atime +nnn] [-mtime +nnn] -print > filename &
```

directory Identifies the directory you want to search. Directories below this directory are also searched.

`-atime +nnn` Finds files that have not been accessed within the number of days (*nnn*) that you specify.

`-mtime +nnn` Finds files that have not been modified within the number of days (*nnn*) that you specify.

filename Identifies the file that contains the list of inactive files.

3 Remove the inactive files found listed in the previous step.

```
# rm 'cat filename'
```

where *filename* identifies the file that was created in the previous step. This file contains the list of inactive files.

Example 8–11 Finding and Removing Old or Inactive Files

The following example shows files in the `/var/adm` directory and the subdirectories that have not been accessed in the last 60 days. The `/var/tmp/deadfiles` file contains the list of inactive files. The `rm` command removes these inactive files.

```
# find /var/adm -type f -atime +60 -print > /var/tmp/deadfiles &
# more /var/tmp/deadfiles
/var/adm/aculog
/var/adm/spellhist
/var/adm/wtmpx
/var/adm/sa/sa13
/var/adm/sa/sa27
/var/adm/sa/sa11
/var/adm/sa/sa23
/var/adm/sulog
/var/adm/vold.log
/var/adm/messages.1
/var/adm/messages.2
/var/adm/messages.3
# rm 'cat /var/tmp/deadfiles'
#
```

▼ How to Clear Out Temporary Directories

1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Change to the directory that you want to clean out.

```
# cd directory
```



Caution – Ensure that you are in the correct directory before completing Step 3. Step 3 deletes all files in the current directory.

3 Delete the files and subdirectories in the current directory.

```
# rm -r *
```

4 Change to other directories that contain unnecessary, temporary or obsolete subdirectories and files.

5 Delete these subdirectories and files by repeating Step 3.

Example 8–12 Clearing Out Temporary Directories

The following example shows how to clear out the `mywork` directory, and how to verify that all files and subdirectories were removed.

```
# cd mywork
# ls
filea.000
fileb.000
filec.001
# rm -r *
# ls
#
```

▼ How to Find and Delete core Files

1 Become an administrator.

Change to the directory where you want to search for core files.

2 Find and remove any core files in this directory and its subdirectories.

```
# find . -name core -exec rm {} \;
```

Example 8–13 Finding and Deleting core Files

The following example shows how to find and remove core files from the `jones` user account by using the `find` command.

```
# cd /home/jones
# find . -name core -exec rm {} \;
```

▼ How to Delete Crash Dump Files

Crash dump files can be very large. If you have enabled your system to store these files, do not retain them for longer than necessary.

1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

2 Change to the directory where crash dump files are stored.

```
# cd /var/crash/system
```

where `system` identifies a system that created the crash dump files.



Caution – Ensure you are in the correct directory before completing Step 3. Step 3 deletes all files in the current directory.

3 Remove the crash dump files.

```
# rm *
```

4 Verify that the crash dump files were removed.

```
# ls
```

Example 8–14 Deleting Crash Dump Files

The following example shows how to remove crash dump files from the system `venus`, and how to verify that the crash dump files were removed.

```
# cd /var/crash/venus  
# rm *  
# ls
```

Administering Disks (Tasks)

This chapter contains disk administration procedures. Many procedures described in this chapter are optional if you are already familiar with how disks are managed on systems running Oracle Solaris

This is a list of the information in this chapter:

- “Administering Disks (Task Map)” on page 189
- “Identifying Disks on a System” on page 190
- “Formatting a Disk” on page 192
- “Displaying Disk Slices” on page 194
- “Creating and Examining a Disk Label” on page 197
- “Recovering a Corrupted Disk Label” on page 203
- “Adding a Third-Party Disk” on page 206

For overview information about disk management, see [Chapter 7, “Managing Disks \(Overview\)”](#).

Administering Disks (Task Map)

Task	Description	For Instructions
Identify the disks on a system.	If you are not sure of the types of disks on a system, use the <code>format</code> utility to identify the disk types.	“How to Identify the Disks on a System” on page 190
Format the disk.	Determine whether a disk is already formatted by using the <code>format</code> utility.	“How to Determine if a Disk Is Formatted” on page 192

Task	Description	For Instructions
	In most cases, disks are already formatted. Use the format utility if you need to format a disk.	“How to Format a Disk” on page 193
Display slice information.	Display slice information by using the format utility.	“How to Display Disk Slice Information” on page 195
Label the disk.	Create the disk label by using the format utility.	“How to Label a Disk” on page 197
Examine the disk label.	Examine the disk label by using the prtvtoc command.	“How to Examine a Disk Label” on page 202
Recover a corrupted disk label.	You can attempt to recover a disk label that was damaged due to a system or power failure.	“How to Recover a Corrupted Disk Label” on page 204

Identifying Disks on a System

Use the format utility to discover the types of disks that are connected to a system. You can also use the format utility to verify that a disk is known to the system. For detailed information on using the format utility, see [Chapter 13, “The format Utility \(Reference\)”](#).

▼ How to Identify the Disks on a System

1 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights” in Oracle Solaris 11.1 Administration: Security Services](#).

2 Identify the disks that are recognized on the system by using the format utility.

```
# format
```

The format utility displays a list of disks that it recognizes under AVAILABLE DISK SELECTIONS.

Example 9-1 Identifying the Disks on a System

The following example shows format command output.

```
# format
AVAILABLE DISK SELECTIONS:
 0. c2t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
   /pci@1c,600000/scsi@2/sd@0,0
 1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
   /pci@1c,600000/scsi@2/sd@1,0
```

- ```

2. c2t2d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
 /pci@1c,600000/scsi@2/sd@2,0
3. c2t3d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
 /pci@1c,600000/scsi@2/sd@3,0

```

Specify disk (enter its number):

The output associates a disk's physical and logical device name to the disk's marketing name, which appears in angle brackets <>. See the example below. This method is an easy way to identify which logical device names represent the disks that are connected to your system. For a description of logical and physical device names, see [“Device Naming Conventions” on page 61](#).

The following example uses a wildcard to display the four disks that are connected to a controller 0:

```

format /dev/rdisk/c0t6*
AVAILABLE DISK SELECTIONS:
 0. /dev/rdisk/c0t600A0B800022024E000054AC4970A629d0p0 <...>
 /scsi_vhci/disk@g600a0b800022024e000054ac4970a629
 1. /dev/rdisk/c0t600A0B800022024E000054AE4970A711d0p0 <...>
 /scsi_vhci/disk@g600a0b800022024e000054ae4970a711
 2. /dev/rdisk/c0t600A0B800022028A000050444970A834d0p0 <....>
 /scsi_vhci/disk@g600a0b800022028a000050444970a834
 3. /dev/rdisk/c0t600A0B800022028A000050454970A8EAd0p0 <...>
 /scsi_vhci/disk@g600a0b800022028a000050454970a8ea
Specify disk (enter its number):

```

The following example shows how to identify the disk information:

```

format
0. c2t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
 /pci@1c,600000/scsi@2/sd@0,0

```

The output identifies that disk 0 (target 0) is connected to the second SCSI host adapter (scsi@2), which is connected to the second PCI interface (/pci@1c,600000/. . .). The output also associates both the physical and logical device name to the disk's marketing name, SUN36G.

Some disks do not have a marketing name. If the format output does not identify disks by their marketing names, then you can use the format utility's type and label features as described in the following steps to include the disk's marketing name.

The following steps must be done while the disk or system is inactive, which means booting from an installation DVD or the network, unless the disk is currently unused and it will not contain the Oracle Solaris release. In addition, the final step is to relabel the disk, which can remove any existing partition information or data.

```

format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
 0. c2t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
 /pci@1c,600000/scsi@2/sd@0,0
 1. c2t1d0 <SEAGATE-ST336607LSUN36G-0307-33.92GB>

```

```

 /pci@1c,600000/scsi@2/sd@1,0
 2. c2t2d0 <SEAGATE-ST336607LSUN36G-0507-33.92GB>
 /pci@1c,600000/scsi@2/sd@2,0
 3. c2t3d0 <drive type unknown>
 /pci@1c,600000/scsi@2/sd@3,0
Specify disk (enter its number): 3
selecting c2t3d0
[disk formatted]
format> type
AVAILABLE DRIVE TYPES:
 0. Auto configure
 1. other
Specify disk type (enter its number): 0
c2t3d0: configured with capacity of 33.92GB
<SEAGATE-ST336607LSUN36G-0507-33.92GB>
[disk formatted]
format> label
Ready to label disk, continue? yes
format> quit

```

**More Information** If the format Utility Does Not Recognize a Disk ...

- Go to [Chapter 10, “Setting Up Disks \(Tasks\),”](#) or [“x86: Setting Up Disks for ZFS File Systems \(Task Map\)”](#) on page 215.
- Go to [“How to Label a Disk”](#) on page 197.
- Connect the disk to the system by using your disk hardware documentation.

## Formatting a Disk

Disks are typically formatted by the manufacturer or reseller. They usually do not need to be reformatted when you install the drive.

A disk must be formatted before you can do the following:

- Write data to the disk. However, most disks are already formatted.
- Use the Oracle Solaris installation utility to install the system.




---

**Caution** – Formatting a disk is a destructive process because it overwrites data on the disk. For this reason, disks are usually formatted only by the manufacturer or reseller. If you think disk defects are the cause of recurring problems, you can use the `format` utility to do a surface analysis. However, be careful to use only the commands that do not destroy data.

---

### ▼ How to Determine if a Disk Is Formatted

**1 Become an administrator.**

For more information, see [“How to Use Your Assigned Administrative Rights”](#) in *Oracle Solaris 11.1 Administration: Security Services*.

**2 Invoke the format utility.**

```
format
```

A numbered list of disks is displayed.

**3 Type the number of the disk that you want to check.**

```
Specify disk (enter its number): 0
```

**4 Verify that the disk you chose is formatted by noting the following message:**

```
[disk formatted]
```

**Example 9–2 Determining if a Disk Is Formatted**

The following example shows that disk `c2t3d0` is formatted.

```
format
AVAILABLE DISK SELECTIONS:
 0. c2t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@1c,600000/scsi@2/sd@0,0
 1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@1c,600000/scsi@2/sd@1,0
 2. c2t2d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
 /pci@1c,600000/scsi@2/sd@2,0
 3. c2t3d0 <SUN146G cyl 14087 alt 2 hd 24 sec 848>
 /pci@1c,600000/scsi@2/sd@3,0
Specify disk (enter its number): 3
selecting c2t3d0
[disk formatted]
```

**▼ How to Format a Disk**

Disks are formatted by the manufacturer. Reformatting a disk should occur rarely. The process is time-consuming and removes all data from the disk.

**1 Become an administrator.**

For more information, see [“How to Use Your Assigned Administrative Rights” in Oracle Solaris 11.1 Administration: Security Services](#).

**2 Invoke the format utility.**

```
format
```

A numbered list of disks is displayed.

**3 Type the number of the disk that you want to format.**

```
Specify disk (enter its number): 3
```




---

**Caution** – Do not select the disk that contains the root file system. If you format a root pool disk, you delete the OS and any data on this disk.

---

- 4 To begin formatting the disk, type `format` at the `format>` prompt. Confirm the command by typing `y`.**

```
format> format
The protection information is not enabled
The disk will be formatted with protection type 0
```

```
Ready to format. Formatting cannot be interrupted
and takes 169 minutes (estimated). Continue? yes
```

- 5 Verify that the disk format was successful by noting the following messages:**

```
Beginning format. The current time is Fri Apr 1 ...
```

```
Formatting...
done
```

```
Verifying media...
 pass 0 - pattern = 0xc6dec6de
14086/23/734
```

```
 pass 1 - pattern = 0x6db6db6d
14086/23/734
```

```
Total of 0 defective blocks repaired.
```

- 6 Exit the `format` utility.**

```
format> quit
```

## Displaying Disk Slices

The best way to create ZFS storage pools is to use whole disks instead of disk slices because whole disks are easier to manage. The only time you need to use a disk slice is when the disk is intended for the ZFS root pool. This is a long-standing boot limitation. For non-root pools, use whole disks. When you create a pool with whole disks, an EFI label is applied. See the EFI disk label example that follows.

If you need to prepare a disk for use as a root pool disk, create a slice 0 that contains the entire disk capacity, as shown in the VTOC disk label example that follows.

For information about setting up disks for use with ZFS storage pools, see [Chapter 10, “Setting Up Disks \(Tasks\)”](#), or [“x86: Setting Up Disks for ZFS File Systems \(Task Map\)”](#) on page 215.

---

**Note** – The `format` utility uses the term *partition* instead of *slice*.

---

## ▼ How to Display Disk Slice Information

You might need to display disk slice information if the disk is intended to be used for the ZFS root pool. It must also include a SMI label.

### 1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Invoke the `format` utility.

```
format
```

A numbered list of disks is displayed.

### 3 Type the number of the disk for which you want to display slice information.

```
Specify disk (enter its number):1
```

### 4 Select the `partition` menu.

```
format> partition
```

### 5 Display the slice information for the selected disk.

```
partition> print
```

### 6 Exit the `format` utility.

```
partition> q
format> q
```

### 7 Verify the displayed slice information by identifying specific slice tags and slices.

If the screen output shows that no slice sizes are assigned, the disk probably does not have slices.

## Example 9-3 Displaying Disk Slice Information

The following example displays slice information for a disk with a VTOC label.

```
format
Searching for disks...done
Specify disk (enter its number):3
Selecting c2t3d0
format> partition
partition> print
Current partition table (c2t3d0):
Total disk cylinders available: 14087 + 2 (reserved cylinders)
```

```

Part Tag Flag Cylinders Size Blocks
 0 root wm 0 - 14086 136.71GB (14087/0/0) 286698624
 1 swap wu 0 0 (0/0/0) 0
 2 backup wu 0 - 14086 136.71GB (14087/0/0) 286698624
 3 unassigned wm 0 0 (0/0/0) 0
 4 unassigned wm 0 0 (0/0/0) 0
 5 unassigned wm 0 0 (0/0/0) 0
 6 usr wm 0 0 (0/0/0) 0
 7 unassigned wm 0 0 (0/0/0) 0
partition> q
format> q

```

For a detailed description of the slice information in these examples, see [Chapter 7, “Managing Disks \(Overview\).”](#)

The following example shows the slice information for a disk with an EFI label:

```

format
Searching for disks...done
Specify disk (enter its number): 3
selecting c2t3d0
[disk formatted]
format> partition
partition> print
Current partition table (default):
Total disk sectors available: 286722878 + 16384 (reserved sectors)

Part Tag Flag First Sector Size Last Sector
 0 usr wm 34 136.72GB 286722911
 1 unassigned wm 0 0 0
 2 unassigned wm 0 0 0
 3 unassigned wm 0 0 0
 4 unassigned wm 0 0 0
 5 unassigned wm 0 0 0
 6 unassigned wm 0 0 0
 7 unassigned wm 0 0 0
 8 reserved wm 286722912 8.00MB 286739295
partition> q
format> q

```

The following example shows the slice information for an EFI (GPT) labeled disk for a root pool:

```

format
Searching for disks...done
Specify disk (enter its number): 0
selecting c2t0d0
[disk formatted]
format> partition
partition> print
Current partition table (original):
Total disk sectors available: 27246525 + 16384 (reserved sectors)
Part Tag Flag First Sector Size Last Sector
 0 BIOS_boot wm 256 256.00MB 524543
 1 usr wm 524544 12.74GB 27246558

```

```

2 unassigned wm 0 0 0
3 unassigned wm 0 0 0
4 unassigned wm 0 0 0
5 unassigned wm 0 0 0
6 unassigned wm 0 0 0
8 reserved wm 27246559 8.00MB 27262942
partition> q
format> q

```

## Creating and Examining a Disk Label

The labeling of a disk is usually done during system installation or when you are using new disks. You might need to relabel a disk if the disk label becomes corrupted. For example, from a power failure.

The `format` utility attempts to automatically configure any unlabeled SCSI disk. If the `format` utility is able to automatically configure an unlabeled disk, it displays a message similar to the following:

```
c2t3d0: configured with capacity of 136.73GB
```

### ▼ How to Label a Disk

If you want to put an EFI label on disk smaller than 2 terabytes, see [Example 9–5](#).

#### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

#### 2 Invoke the `format` utility.

```
format
```

A numbered list of disks is displayed.

#### 3 Type the number of the disk that you want to label.

```
Specify disk (enter its number):1
```

If the `format` utility recognizes the disk type, the next step is to search for a backup label to label the disk. Labeling the disk with the backup label labels the disk with the correct partitioning information, the disk type, and disk geometry.

#### 4 Select one of the following to label the disk:

- If the disk is unlabeled and was successfully configured, go to Step 5 to label the disk.

The `format` utility will ask if you want to label the disk.

- If the disk is labeled but you want to change the disk type, or if the format utility was not able to automatically configure the disk, proceed to Step 6 to set the disk type and label the disk.

**5 Label the disk by typing y at the Label it now? prompt.**

Disk not labeled. Label it now? **y**

The disk is now labeled. Go to step 10 to exit the format utility.

**6 Enter type at the format> prompt.**

format> **type**

The Available Drive Types menu is displayed.

**7 Select a disk type from the list of possible disk types.**

Specify disk type (enter its number)[12]: **12**

Or, select 0 to automatically configure a SCSI-2 disk.

**8 Label the disk. If the disk is not labeled, the following message is displayed.**

Disk not labeled. Label it now? **y**

Otherwise, you are prompted with this message:

Ready to label disk, continue? **y**

**9 Verify the disk label.**

format> **verify**

**10 Exit the format utility.**

format> **q**

#

**Example 9-4 Labeling a Disk**

The following example shows how to automatically configure and label a 36-GB disk.

```
format
AVAILABLE DISK SELECTIONS:
 0. c0t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
 /pci@1c,600000/scsi@2/sd@0,0
 1. c0t1d0 <SEAGATE-ST336607LSUN36G-0307-33.92GB>
 /pci@1c,600000/scsi@2/sd@1,0
 2. c0t2d0 <SEAGATE-ST336607LSUN36G-0507-33.92GB>
 /pci@1c,600000/scsi@2/sd@2,0
 3. c0t3d0 <SEAGATE-ST336607LSUN36G-0507-33.92GB>
 /pci@1c,600000/scsi@2/sd@3,0
Specify disk (enter its number): 3
c0t3d0: configured with capacity of 33.92GB
Disk not labeled. Label it now? yes
```

```
format> verify
format> q
```

### Example 9-5 Labeling a Disk With an EFI Label

The following example shows how to use the `format -e` command to label a disk with an EFI label or an SMI (VTOC) label. Remember to verify that your layered software products will continue to work on systems with EFI-labeled disks. For general information on EFI label restrictions, see [“Restrictions of the EFI Disk Label” on page 165](#).

```
format -e
Searching for disks...done
AVAILABLE DISK SELECTIONS:
 0. c2t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@1c,600000/scsi@2/sd@0,0
 1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@1c,600000/scsi@2/sd@1,0
 2. c2t2d0 <FUJITSU-MAP3147N SUN146G-0501-136.73GB>
 /pci@1c,600000/scsi@2/sd@2,0
 3. c2t3d0 <FUJITSU-MAP3147N SUN146G-0501-136.73GB>
 /pci@1c,600000/scsi@2/sd@3,0
Specify disk (enter its number): 3
selecting c2t3d0
[disk formatted]
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[0]: 1
Ready to label disk, continue? yes
format> quit
```

### Example 9-6 Change a EFI-Labeled Disk to an SMI-Labeled Disk

The following example shows how to use the `format -e` command to change an EFI labeled disk to an SMI-labeled disk that can be used for a ZFS root pool.

On an x86 based system, you will first have to change the EFI fdisk partition to a Solaris partition. For example:

```
format -e
select disk ...
format> fdisk
FORMAT MENU:
 disk - select a disk
 type - select (define) a disk type
 Total disk size is 17833 cylinders
 Cylinder size is 16065 (512 byte) blocks

Partition Status Type Start End Length %
===== ===== =====
 1 EFI 0 17833 17834 100
```

```

SELECT ONE OF THE FOLLOWING:
 1. Create a partition
 2. Specify the active partition
 3. Delete a partition
 4. Change between Solaris and Solaris2 Partition IDs
 5. Edit/View extended partitions
 6. Exit (update disk configuration and exit)
 7. Cancel (exit without updating disk configuration)
Enter Selection: 3
Specify the partition number to delete (or enter 0 to exit): 1
This will make all files and
programs in this partition inaccessible (type "y" or "n"). y
Enter Selection: 1
Select the partition type to create:
 1=SOLARIS2 2=UNIX 3=PCIX05 4=Other 5=DOS12
 6=DOS16 7=DOSEXT 8=DOSBIG 9=DOS16LBA A=x86 Boot
 B=Diagnostic C=FAT32 D=FAT32LBA E=DOSEXTLBA F=EFI (Protective)
 G=EFI SYS 0=Exit? 1
Specify the percentage of disk to use for this partition
(or type "c" to specify the size in cylinders). 100
Should this become the active partition? If yes, it will be activated
each time the computer is reset or turned on.
Please type "y" or "n". y
Enter Selection: 6
Partition 1 is now the active partition.

```

On a SPARC based system, follow these steps:

```

format -e
Searching for disks...done
AVAILABLE DISK SELECTIONS:
 0. c2t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@1c,600000/scsi@2/sd@0,0
 1. c2t1d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@1c,600000/scsi@2/sd@1,0
 2. c2t2d0 <FUJITSU-MAP3147N SUN146G-0501-136.73GB>
 /pci@1c,600000/scsi@2/sd@2,0
 3. c2t3d0 <FUJITSU-MAP3147N SUN146G-0501-136.73GB>
 /pci@1c,600000/scsi@2/sd@3,0
Specify disk (enter its number): 3
selecting c2t0d0
[disk formatted]
format> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 0
Ready to label disk, continue? yes

```

You should also check the default partition table to ensure that it is optimal for a root pool slice, which means that the bulk of the disk space is in slice 0. See the steps below to increase the size of slice 0.

```

format> partition
partition> print
Current partition table (default):

```

Total disk cylinders available: 14085 + 2 (reserved cylinders)

| Part | Tag        | Flag | Cylinders  | Size     | Blocks                |
|------|------------|------|------------|----------|-----------------------|
| 0    | root       | wm   | 1 - 13     | 129.19MB | (13/0/0) 264576       |
| 1    | swap       | wu   | 14 - 26    | 129.19MB | (13/0/0) 264576       |
| 2    | backup     | wu   | 0 - 14086  | 136.71GB | (14087/0/0) 286698624 |
| 3    | unassigned | wm   | 0          | 0        | (0/0/0) 0             |
| 4    | unassigned | wm   | 0          | 0        | (0/0/0) 0             |
| 5    | unassigned | wm   | 0          | 0        | (0/0/0) 0             |
| 6    | usr        | wm   | 27 - 14084 | 136.43GB | (14058/0/0) 286108416 |
| 7    | unassigned | wm   | 0          | 0        | (0/0/0) 0             |
| 8    | boot       | wu   | 0 - 0      | 9.94MB   | (1/0/0) 20352         |
| 9    | alternates | wm   | 0          | 0        | (0/0/0) 0             |

partition> **modify**

Select partitioning base:

0. Current partition table (default)

1. All Free Hog

Choose base (enter number) [0]? **1**

| Part | Tag        | Flag | Cylinders | Size     | Blocks                |
|------|------------|------|-----------|----------|-----------------------|
| 0    | root       | wm   | 0         | 0        | (0/0/0) 0             |
| 1    | swap       | wu   | 0         | 0        | (0/0/0) 0             |
| 2    | backup     | wu   | 0 - 14084 | 136.69GB | (14085/0/0) 286657920 |
| 3    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 4    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 5    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 6    | usr        | wm   | 0         | 0        | (0/0/0) 0             |
| 7    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 8    | boot       | wu   | 0 - 0     | 9.94MB   | (1/0/0) 20352         |
| 9    | alternates | wm   | 0         | 0        | (0/0/0) 0             |

Do you wish to continue creating a new partition

table based on above table[yes]? **yes**

Free Hog partition[6]? **0**

Enter size of partition '1' [0b, 0c, 0.00mb, 0.00gb]:

Enter size of partition '3' [0b, 0c, 0.00mb, 0.00gb]:

Enter size of partition '4' [0b, 0c, 0.00mb, 0.00gb]:

Enter size of partition '5' [0b, 0c, 0.00mb, 0.00gb]:

Enter size of partition '6' [0b, 0c, 0.00mb, 0.00gb]:

Enter size of partition '7' [0b, 0c, 0.00mb, 0.00gb]:

| Part | Tag        | Flag | Cylinders | Size     | Blocks                |
|------|------------|------|-----------|----------|-----------------------|
| 0    | root       | wm   | 1 - 14084 | 136.68GB | (14084/0/0) 286637568 |
| 1    | swap       | wu   | 0         | 0        | (0/0/0) 0             |
| 2    | backup     | wu   | 0 - 14084 | 136.69GB | (14085/0/0) 286657920 |
| 3    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 4    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 5    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 6    | usr        | wm   | 0         | 0        | (0/0/0) 0             |
| 7    | unassigned | wm   | 0         | 0        | (0/0/0) 0             |
| 8    | boot       | wu   | 0 - 0     | 9.94MB   | (1/0/0) 20352         |
| 9    | alternates | wm   | 0         | 0        | (0/0/0) 0             |

Okay to make this the current partition table[yes]? **yes**

Enter table name (remember quotes): **"c2t0d0"**

Ready to label disk, continue? **yes**

partition> **quit**

format> **quit**

## ▼ How to Examine a Disk Label

Examine disk label information by using the `prtvtoc` command. For a detailed description of the disk label and the information that is displayed by the `prtvtoc` command, see [Chapter 7, “Managing Disks \(Overview\)”](#).

### 1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Display the disk label information.

```
prtvtoc /dev/rdisk/device-name
```

where *device-name* is the raw disk device you want to examine.

## Example 9-7 Examining a Disk Label

The following example shows disk label information for a disk with a VTOC label:

```
prtvtoc /dev/rdisk/c2t3d0s0
* /dev/rdisk/c2t3d0s0 partition map
*
* Dimensions:
* 512 bytes/sector
* 848 sectors/track
* 24 tracks/cylinder
* 20352 sectors/cylinder
* 14089 cylinders
* 14087 accessible cylinders
*
* Flags:
* 1: unmountable
* 10: read-only
*
*
* Partition Tag Flags First Sector Sector Last
* 0 2 00 0 286698624 286698623
* 2 5 01 0 286698624 286698623
```

The following example shows disk label information for a root pool disk with an EFI label:

```
prtvtoc /dev/dsk/c7t0d0s1
* /dev/dsk/c7t0d0s1 partition map
*
* Dimensions:
* 512 bytes/sector
* 156301488 sectors
* 156301421 accessible sectors
*
* Flags:
* 1: unmountable
* 10: read-only
```

```

*
* Unallocated space:
* First Sector Last
* Sector Count Sector
* 34 222 255
*
*
* Partition Tag Flags First Sector Last
* Partition Tag Flags Sector Count Sector Mount Directory
* 0 24 00 256 524288 524543
* 1 4 00 524544 155760527 156285070
* 8 11 00 156285071 16384 156301454
~#

```

The following example shows disk label information for a non-root pool disk with an EFI label:

```

prtvtoc /dev/dsk/c8t3d0
* /dev/dsk/c8t3d0 partition map
*
* Dimensions:
* 512 bytes/sector
* 143374738 sectors
* 143374671 accessible sectors
*
* Flags:
* 1: unmountable
* 10: read-only
*
* Unallocated space:
* First Sector Last
* Sector Count Sector
* 34 222 255
*
*
* Partition Tag Flags First Sector Last
* Partition Tag Flags Sector Count Sector Mount Directory
* 0 4 00 256 143358065 143358320
* 8 11 00 143358321 16384 143374704

```

## Recovering a Corrupted Disk Label

Sometimes, a power or system failure causes a disk's label to become unrecognizable. A corrupted disk label doesn't always mean that the slice information or the disk's data must be re-created or restored.

The first step to recovering a corrupted disk label is to label the disk with the correct geometry and disk type information. You can complete this step through the normal disk labeling method, by using either automatic configuration or manual disk type specification.

If the format utility recognizes the disk type, the next step is to search for a backup label to label the disk. Labeling the disk with the backup label labels the disk with the correct partitioning information, the disk type, and disk geometry.

## ▼ How to Recover a Corrupted Disk Label

### 1 Boot the system to single-user mode.

If necessary, boot the system from a local Oracle Solaris DVD or the network in single-user mode to access the disk.

For information on booting the system, see *Booting and Shutting Down Oracle Solaris 11.1 Systems*.

### 2 Relabel the disk.

```
format
```

The format utility attempts to automatically configure any unlabeled SCSI disk. If the format utility is able to configure the unlabeled and corrupted disk, it will display this message:

```
cwtxdy: configured with capacity of abcMB
```

The format utility then displays a numbered list of disks on the system.

### 3 Type the number of the disk that you need to recover.

```
Specify disk (enter its number): 1
```

### 4 Select one of the following to determine how to label the disk.

- If the disk was configured successfully, do the following:

#### a. Search for the backup label.

```
format> verify
Warning: Could not read primary label.
Warning: Check the current partitioning and 'label' the disk or
use the 'backup' command.
Backup label contents:
Volume name = < >
ascii name = <FUJITSU-MAP3147N SUN146G-0501>
pcyl = 14089
ncyl = 14087
acyl = 2
nhead = 24
nsect = 848
Part Tag Flag Cylinders Size Blocks
 0 root wm 0 - 14086 136.71GB (14087/0/0) 286698624
 1 swap wu 0 0 (0/0/0) 0
 2 backup wu 0 - 14086 136.71GB (14087/0/0) 286698624
 3 unassigned wm 0 0 (0/0/0) 0
 4 unassigned wm 0 0 (0/0/0) 0
 5 unassigned wm 0 0 (0/0/0) 0
 6 usr wm 0 0 (0/0/0) 0
 7 unassigned wm 0 0 (0/0/0) 0
```

- b. If the format utility was able to find a backup label, and the backup label contents appear satisfactory, label the disk with the backup label.**

```
format> backup
Disk has a primary label, still continue? y

Searching for backup labels...found.
Restoring primary label

The disk label has been recovered
```

- c. Exit the format utility.**

```
format> q
```

- **If the disk was *not* configured successfully, do the following:**

- a. Specify the disk type by using the type command:**

```
format> type

The Available Drives Type menu is displayed.
```

- b. Select 0 to automatically configure the disk. Or, select a disk type from the list of possible disk types.**

```
Specify disk type (enter its number)[12]: 12
```

- c. If the disk was successfully configured, reply with no when the format utility asks if you want to label the disk.**

```
Disk not labeled. Label it now? no
```

- d. Search for the backup label.**

```
format> verify
Warning: Could not read primary label.
Warning: Check the current partitioning and 'label' the disk
or use the 'backup' command.
.
.
.
```

- e. If the format utility was able to find a backup label, and the backup label contents appear satisfactory, label the disk with the backup label.**

```
format> backup
Disk has a primary label, still continue? y
Searching for backup labels...found.
Restoring primary label

The disk label has been recovered.
```

- f. Exit the format utility.**

```
format> q
```

## 5 Verify the file systems on the recovered disk.

For information on using the `zpool scrub` command for ZFS file systems, see [zpool\(1M\)](#).

# Adding a Third-Party Disk

Oracle Solaris supports many third-party disks. However, for the disk to be recognized, you might need to supply a device driver.

Other options for adding disks are as follows:

- If you are adding a SCSI disk, you might try the `format` utility's automatic configuration feature.
- You might try hot-plugging a PCI, SCSI, or USB disk. For more information, see [Chapter 3, “Managing Devices \(Tasks\)”](#).

---

**Note** – Oracle cannot guarantee that its `format` utility will work properly with all third-party disk drivers. If the disk driver is not compatible with the `format` utility, the disk drive vendor should supply you with a custom disk formatting program.

---

Typically, you discover that software support is missing when you invoke the `format` utility and find that the disk type is not recognized.

Refer to the appropriate configuration procedure for adding system disks or secondary disks in [Chapter 10, “Setting Up Disks \(Tasks\)”](#), or [“x86: Setting Up Disks for ZFS File Systems \(Task Map\)”](#) on page 215.

# Setting Up Disks (Tasks)

---

This chapter describes how to set up disks.

This is a list of the information in this chapter:

- “SPARC: Setting up Disks (Task Map)” on page 207
- “x86: Setting Up Disks for ZFS File Systems (Task Map)” on page 215
- “x86: Creating and Changing Solaris `fdisk` Partitions” on page 229

For overview information about disk management, see Chapter 7, “Managing Disks (Overview).”

For step-by-step instructions on setting up disks on an x86 based system, see “x86: Setting Up Disks for ZFS File Systems (Task Map)” on page 215.

## SPARC: Setting up Disks (Task Map)

The following task map identifies the procedures for setting up a ZFS root pool disk for a ZFS root file system or a non-root ZFS pool disk on a SPARC based system.

| Task                                           | Description                                                                                                                                               | For Instructions                                                     |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| 1. Set up the disk for a ZFS root file system. | <i>Disk for a ZFS Root File System</i><br>Connect the new disk or replace the existing root pool disk and boot from a local or remote Oracle Solaris DVD. | “SPARC: How to Set Up a Disk for a ZFS Root File System” on page 209 |

| Task                                                                 | Description                                                                                                                                                                                                     | For Instructions                                                                           |
|----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| 2. Install the boot blocks for a ZFS root file system, if necessary. | If you replace a disk that is intended for the root pool by using the <code>zpool replace</code> command, then you must install the boot blocks manually so that the system can boot from the replacement disk. | <a href="#">“SPARC: How to Install Boot Blocks for a ZFS Root File System” on page 214</a> |
| 3. Set up a disk for ZFS non-root file system.                       | <i>Disk for a ZFS Non-Root File System</i><br>Set up a disk for a ZFS non-root file system.                                                                                                                     | <a href="#">“SPARC: How to Set Up a Disk for a ZFS Non-Root File System” on page 214</a>   |

## SPARC: Setting Up Disks for ZFS File Systems

Although the procedures that describe how to set up a disk can be used with a ZFS file system, a ZFS file system is not directly mapped to a disk or a disk slice. You must create a ZFS storage pool before creating a ZFS file system. For more information, see [Oracle Solaris 11.1 Administration: ZFS File Systems](#).

The root pool contains the root file system that is used to boot the Oracle Solaris OS. If a root pool disk becomes damaged and the root pool is not mirrored, the system might not boot.

If a root pool disk becomes damaged, you have two ways to recover:

- You can reinstall the entire Oracle Solaris OS.
- Or, you can replace the root pool disk and restore your file systems from snapshots or from a backup medium. You can reduce system down time due to hardware failures by creating a redundant root pool. The only supported redundant root pool configuration is a mirrored root pool.

A disk that is used in a non-root pool usually contains space for user or data files. You can attach or add another disk to a root pool or a non-root pool for more disk space.

Or, you can replace a damaged disk in a pool in the following ways:

- A disk can be replaced in a non-redundant pool, if all of the devices are currently ONLINE.
- A disk can be replaced in a redundant pool, if enough redundancy exists among the other devices.
- In a mirrored root pool, you can replace a disk or attach a disk and then detach the failed disk or a smaller disk to increase a pool's size.

In general, setting up a disk on the system depends on the hardware, so review your hardware documentation when adding or replacing a disk on your system. If you need to add a disk to an

existing controller, then it might just be a matter of inserting the disk in an empty slot, if the system supports hot-plugging. If you need to configure a new controller, see [“Dynamic Reconfiguration and Hot-Plugging”](#) on page 65.

## ▼ SPARC: How to Set Up a Disk for a ZFS Root File System

Refer to your hardware installation guide for information on replacing a disk.

- 1 **Disconnect the damaged disk from the system, if necessary.**
- 2 **Connect the replacement disk to the system and check the disk's physical connections, if necessary.**
- 3 **Follow the instructions in the following table, depending on whether you are booting from a local Oracle Solaris DVD or a remote Oracle Solaris DVD from the network.**

| Boot Type                                   | Action                                                                                                                         |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| From an Oracle Solaris DVD in a local drive | 1. Make sure the Oracle Solaris DVD is in the drive.<br>2. Boot from the media to single-user mode:<br>ok <b>boot cdrom -s</b> |
| From the network                            | Boot from the network to single-user mode:<br>ok <b>boot net:dhcp</b>                                                          |

After a few minutes, select option 3 - Shell.

### More Information After You Set Up a Disk for a ZFS Root File System ...

After the disk is connected or replaced, you can create a slice and update the disk label. Go to [“SPARC: How to Create a Disk Slice for a ZFS Root File System”](#) on page 210.

## SPARC: Creating a Disk Slice for a ZFS Root File System

You must create a disk slice for a disk that is intended for a ZFS root pool on SPARC systems that do not have GPT-aware firmware. This is a long-standing boot limitation.

Review the following root pool disk requirements:

- In Oracle Solaris 11.1, an EFI (GPT) label is installed on a SPARC system with GPT aware firmware and on an x86 system. Otherwise, an SMI (VTOC) label is installed.
- Must be a single disk or be part of a mirrored configuration. Neither a non-redundant configuration nor a RAIDZ configuration is supported for the root pool.
- All subdirectories of the root file system that are part of the OS image, with the exception of `/var`, must be in the same dataset as the root file system.
- All Oracle Solaris OS components must reside in the root pool, with the exception of the swap and dump devices.
- For a root pool disk that is labeled with VTOC, you should create a disk slice with the bulk of disk space in slice 0, if you need to replace a root pool disk.
- Attempting to use different slices on a disk and share that disk among different operating systems or with a different ZFS storage pool or storage pool components is not recommended.

## ▼ SPARC: How to Create a Disk Slice for a ZFS Root File System

In general, the root pool disk is installed automatically when the system is installed. If you need to replace a root pool disk or attach a new disk as a mirrored root pool disk, see the steps that follow.

### 1 Become an administrator.

### 2 Offline and unconfigure the failed disk, if necessary.

Some hardware requires that you offline and unconfigure a disk before attempting the `zpool replace` operation to replace a failed disk. For example:

```
zpool offline rpool c2t1d0s0
cfgadm -c unconfigure c2::disk/c2t1d0
```

### 3 Physically connect the new or replacement disk to the system, if necessary.

#### a. Physically remove the failed disk.

#### b. Physically insert the replacement disk.

#### c. Configure the replacement disk, if necessary. For example:

```
cfgadm -c configure c2::disk/c2t1d0
```

On some hardware, you do not have to reconfigure the replacement disk after it is inserted.

#### 4 Confirm that the disk is accessible by reviewing the format output.

For example, the format command shows 4 disks connected to this system.

```
format -e
AVAILABLE DISK SELECTIONS:
 0. c2t0d0 <SUN36G cyl 24620 alt 2 hd 27 sec 107>
 /pci@1c,600000/scsi@2/sd@0,0
 1. c2t1d0 <SEAGATE-ST336607LSUN36G-0307-33.92GB>
 /pci@1c,600000/scsi@2/sd@1,0
 2. c2t2d0 <SEAGATE-ST336607LSUN36G-0507-33.92GB>
 /pci@1c,600000/scsi@2/sd@2,0
 3. c2t3d0 <SEAGATE-ST336607LSUN36G-0507-33.92GB>
 /pci@1c,600000/scsi@2/sd@3,0
```

#### 5 Select the disk to be used for the ZFS root pool.

#### 6 Confirm that the disk has an SMI label by displaying the partition (slice) information.

For example, the partition (slice) output for c2t1d0 shows that this disk has an EFI label because it identifies first and last sectors.

```
Specify disk (enter its number): 1
selecting c2t1d0
[disk formatted]
format> p
PARTITION MENU:
 0 - change '0' partition
 1 - change '1' partition
 2 - change '2' partition
 3 - change '3' partition
 4 - change '4' partition
 5 - change '5' partition
 6 - change '6' partition
expand - expand label to use whole disk
select - select a predefined table
modify - modify a predefined partition table
name - name the current table
print - display the current table
label - write partition map and label to the disk
!<cmd> - execute <cmd>, then return
quit
partition> p
Current partition table (original):
Total disk sectors available: 71116508 + 16384 (reserved sectors)

Part Tag Flag First Sector Size Last Sector
 0 usr wm 256 33.91GB 71116541
 1 unassigned wm 0 0 0
 2 unassigned wm 0 0 0
 3 unassigned wm 0 0 0
 4 unassigned wm 0 0 0
 5 unassigned wm 0 0 0
 6 unassigned wm 0 0 0
 8 reserved wm 71116542 8.00MB 71132925

partition>
```

## 7 If the disk contains an EFI label, relabel the disk with an SMI label.

For example, the `c2t1d0` disk is relabeled with an SMI label, but the default partition table does not provide an optimal slice configuration.

```
partition> label
[0] SMI Label
[1] EFI Label
Specify Label type[1]: 0
Auto configuration via format.dat[no]?
Auto configuration via generic SCSI-2[no]?
partition> p
Current partition table (default):
Total disk cylinders available: 24620 + 2 (reserved cylinders)
```

| Part | Tag        | Flag | Cylinders   | Size     | Blocks               |
|------|------------|------|-------------|----------|----------------------|
| 0    | root       | wm   | 0 - 90      | 128.37MB | (91/0/0) 262899      |
| 1    | swap       | wu   | 91 - 181    | 128.37MB | (91/0/0) 262899      |
| 2    | backup     | wu   | 0 - 24619   | 33.92GB  | (24620/0/0) 71127180 |
| 3    | unassigned | wm   | 0           | 0        | (0/0/0) 0            |
| 4    | unassigned | wm   | 0           | 0        | (0/0/0) 0            |
| 5    | unassigned | wm   | 0           | 0        | (0/0/0) 0            |
| 6    | usr        | wm   | 182 - 24619 | 33.67GB  | (24438/0/0) 70601382 |
| 7    | unassigned | wm   | 0           | 0        | (0/0/0) 0            |

```
partition>
```

## 8 Create an optimal slice configuration for a ZFS root pool disk.

Set the free hog partition so that all the unallocated disk space is collected in slice 0. Then, press return through the slice size fields to create one large slice 0.

```
partition> modify
Select partitioning base:
 0. Current partition table (default)
 1. All Free Hog
Choose base (enter number) [0]? 1
```

| Part | Tag        | Flag | Cylinders | Size    | Blocks               |
|------|------------|------|-----------|---------|----------------------|
| 0    | root       | wm   | 0         | 0       | (0/0/0) 0            |
| 1    | swap       | wu   | 0         | 0       | (0/0/0) 0            |
| 2    | backup     | wu   | 0 - 24619 | 33.92GB | (24620/0/0) 71127180 |
| 3    | unassigned | wm   | 0         | 0       | (0/0/0) 0            |
| 4    | unassigned | wm   | 0         | 0       | (0/0/0) 0            |
| 5    | unassigned | wm   | 0         | 0       | (0/0/0) 0            |
| 6    | usr        | wm   | 0         | 0       | (0/0/0) 0            |
| 7    | unassigned | wm   | 0         | 0       | (0/0/0) 0            |

```
Do you wish to continue creating a new partition
table based on above table[yes]?
```

```
Free Hog partition[6]? 0
```

```
Enter size of partition '1' [0b, 0c, 0.00mb, 0.00gb]:
```

```
Enter size of partition '3' [0b, 0c, 0.00mb, 0.00gb]:
```

```
Enter size of partition '4' [0b, 0c, 0.00mb, 0.00gb]:
```

```
Enter size of partition '5' [0b, 0c, 0.00mb, 0.00gb]:
```

```
Enter size of partition '6' [0b, 0c, 0.00mb, 0.00gb]:
```

```
Enter size of partition '7' [0b, 0c, 0.00mb, 0.00gb]:
```

| Part | Tag | Flag | Cylinders | Size | Blocks |
|------|-----|------|-----------|------|--------|
|------|-----|------|-----------|------|--------|

```

0 root wm 0 - 24619 33.92GB (24620/0/0) 71127180
1 swap wu 0 0 (0/0/0) 0
2 backup wu 0 - 24619 33.92GB (24620/0/0) 71127180
3 unassigned wm 0 0 (0/0/0) 0
4 unassigned wm 0 0 (0/0/0) 0
5 unassigned wm 0 0 (0/0/0) 0
6 usr wm 0 0 (0/0/0) 0
7 unassigned wm 0 0 (0/0/0) 0

```

```

Okay to make this the current partition table[yes]?
Enter table name (remember quotes): "c2t1d0"

```

```

Ready to label disk, continue? yes
partition> quit
format> quit

```

### 9 Let ZFS know that the failed disk is replaced.

```

zpool replace rpool c2t1d0s0
zpool online rpool c2t1d0s0

```

On some hardware, you do not have to online the replacement disk after it is inserted.

If you are attaching a new disk to create a mirrored root pool or attaching a larger disk to replace a smaller disk, use syntax similar to the following:

```

zpool attach rpool c2t0d0s0 c2t1d0s0

```

A `zpool attach` operation on a root pool disk applies the boot blocks automatically.

### 10 If a root pool disk is replaced with a new disk, apply the boot blocks after the new or replacement disk is resilvered.

For example:

```

zpool status rpool
bootadm install-bootloader

```

A `zpool replace` operation on a root pool disk does not apply the boot blocks automatically.

### 11 Verify that you can boot from the new disk.

### 12 If the system boots from the new disk, detach the old disk.

This step is only necessary if you attach a new disk to replace a failed disk or a smaller disk.

```

zpool detach rpool c2t0d0s0

```

### 13 Set up the system to boot automatically from the new disk, either by using the `eeeprom` command or the `setenv` command from the SPARC boot PROM.

## ▼ SPARC: How to Install Boot Blocks for a ZFS Root File System

- 1 **Become an administrator.**
- 2 **Install a boot block for a ZFS root file system.**  

```
bootadm install-bootloader
```

For more information, see [installboot\(1M\)](#).
- 3 **Verify that the boot blocks are installed by rebooting the system to run level 3.**  

```
init 6
```

### Example 10-1 SPARC: Installing Boot Blocks for a ZFS Root File System

If you physically replace the disk that is intended for the root pool and the Oracle Solaris OS is then reinstalled, or you attach a new disk for the root pool, the boot blocks are installed automatically. If you replace a disk that is intended for the root pool by using the `zpool replace` command, then you must install the boot blocks manually so that the system can boot from the replacement disk.

The following example shows how to install boot blocks for a ZFS root file system.

```
bootadm install-bootloader
```

## ▼ SPARC: How to Set Up a Disk for a ZFS Non-Root File System

If you are setting up a disk to be used with a non-root ZFS file system, the disk is relabeled automatically when the pool is created or when the disk is added to the pool. If a pool is created with whole disks or when a whole disk is added to a ZFS storage pool, an EFI label is applied. For more information about EFI disk labels, see “[EFI \(GPT\) Disk Label](#)” on page 163.

Generally, most modern bus types support hot-plugging. This means you can insert a disk in an empty slot and the system recognizes it. For more information about hot-plugging devices, see [Chapter 4, “Dynamically Configuring Devices \(Tasks\)”](#).

- 1 **Become an administrator.**
- 2 **Connect the disk to the system and check the disk's physical connections.**  
Refer to the disk's hardware installation guide for details.

**3 Offline and unconfigure the failed disk, if necessary.**

Some hardware requires that you offline and unconfigure a disk before attempting the `zpool replace` operation to replace a failed disk. For example:

```
zpool offline tank c1t1d0
cftadm -c unconfigure c1::dsk/c1t1d0
<Physically remove failed disk c1t1d0>
<Physically insert replacement disk c1t1d0>
cftadm -c configure c1::dsk/c1t1d0
```

On some hardware, you do not to reconfigure the replacement disk after it is inserted.

**4 Confirm that the new disk is recognized.**

Review the output of the `format` utility to see if the disk is listed under AVAILABLE DISK SELECTIONS. Then, quit the `format` utility.

```
format
```

**5 Let ZFS know that the failed disk is replaced, if necessary.**

```
zpool replace tank c1t1d0
zpool online tank c1t1d0
```

Confirm that the new disk is resilvering.

```
zpool status tank
```

**6 Attach a new disk to an existing ZFS storage pool, if necessary.**

For example:

```
zpool attach tank mirror c1t0d0 c2t0d0
```

Confirm that the new disk is resilvering.

```
zpool status tank
```

For more information, see [Chapter 3, “Managing Oracle Solaris ZFS Storage Pools,”](#) in *Oracle Solaris 11.1 Administration: ZFS File Systems*.

## x86: Setting Up Disks for ZFS File Systems (Task Map)

The following task map identifies the procedures for setting up a ZFS root pool disk for a ZFS root file system on an x86 based system.

Task	Description	For Instructions
1. Set up the disk for a ZFS root file system.	<i>Disk for a ZFS Root File System</i> Connect the new disk or replace the existing root pool disk and boot from a local or remote Oracle Solaris DVD.	“x86: How to Set Up a Disk for a ZFS Root File System” on page 217
2. Create or change an fdisk partition, if necessary.	The disk must contain a valid Solaris fdisk partition.	“x86: Creating and Changing Solaris fdisk Partitions” on page 229
3. Recreate the root pool or create an alternate root pool.	Recreate the root pool or alternate root pool, in case of a failure.	“x86: How to Create a Solaris fdisk Partition” on page 230 and “x86: How to Replace a ZFS Root Pool Disk (EFI (GPT))” on page 222 or “x86: How to Replace a ZFS Root Pool Disk (VTOC)” on page 224
4. Install the boot loader if you are replacing a root pool disk by using the zpool replace command.	If you replace a disk that is intended for the root pool by using the zpool replace command, then you must install the boot loader manually so that the system can boot from the replacement disk.	“x86: How to Install Boot Blocks for a ZFS Root File System” on page 227
5. Set up a disk for a ZFS non-root file system.	<i>Disk for a ZFS Non-Root File System</i> Connect the disk.	“x86: How to Set Up a Disk for a ZFS Non-Root File System” on page 228

## x86: Setting Up Disks for ZFS File Systems

Although the procedures that describe how to set up a disk and create an fdisk partition can be used with a ZFS file systems, a ZFS file system is not directly mapped to a disk or a disk slice. You must create a ZFS storage pool before creating a ZFS file system. For more information, see [Oracle Solaris 11.1 Administration: ZFS File Systems](#).

The root pool contains the root file system that is used to boot the Oracle Solaris OS. If a root pool disk becomes damaged and the root pool is not mirrored, the system might not boot.

If a root pool disk becomes damaged, you have two ways to recover:

- You can reinstall the entire Oracle Solaris OS.
- Or, you can replace the root pool disk and restore your file systems from snapshots or from a backup medium. You can reduce system down time due to hardware failures by creating a redundant root pool. The only supported redundant root pool configuration is a mirrored root pool.

A disk that is used in a non-root pool usually contains space for user or data files. You can attach or add another disk to a root pool or a non-root pool for more disk space.

Or, you can replace a damaged disk in a pool in the following ways:

- A disk can be replaced in a non-redundant pool if all the devices are currently ONLINE.
- A disk can be replaced in a redundant pool if enough redundancy exists among the other devices.
- In a mirrored root pool, you can replace a disk or attach a disk and then detach the failed disk or a smaller disk to increase a pool's size.

In general, setting up a disk on the system depends on the hardware so review your hardware documentation when adding or replacing a disk on your system. If you need to add a disk to an existing controller, then it might just be a matter of inserting the disk in an empty slot, if the system supports hot-plugging. If you need to configure a new controller, see [“Dynamic Reconfiguration and Hot-Plugging” on page 65](#).

## ▼ x86: How to Set Up a Disk for a ZFS Root File System

Refer to your hardware installation guide for information on replacing a disk.

- 1 **Disconnect the damaged disk from the system, if necessary.**
- 2 **Connect the replacement disk to the system, and check the disk's physical connections.**
- 3 **Follow the instructions in the following table, depending on whether you are booting from a local Oracle Solaris DVD or a remote Oracle Solaris DVD from the network.**

Boot Type	Action
From an Oracle Solaris DVD in a local drive	<ol style="list-style-type: none"> <li>1. Make sure the Oracle Solaris DVD is in the drive.</li> <li>2. Select the option to boot from the media</li> </ol>
From the network	<ol style="list-style-type: none"> <li>3. Select the option to boot from the network.</li> </ol>

## x86: Preparing a Disk for a ZFS Root File System

Review the following root pool disk requirements:

- In most cases, Oracle Solaris 11.1 installs an EFI (GPT) label for the root pool disk or disks. The SMI (VTOC) label is still available and supported. Follow the procedures in this section based on the EFI (GPT) or SMI (VTOC) labeling.
- Must be a single disk or be part of mirrored configuration. Neither a non-redundant configuration nor a RAIDZ configuration is supported for the root pool.

- All subdirectories of the root file system that are part of the OS image, with the exception of `/var`, must be in the same dataset as the root file system.
- All Oracle Solaris OS components must reside in the root pool, with the exception of the swap and dump devices.
- For x86 systems with a root pool disk that is labeled with EFI, then the correct boot partitions are created automatically, in most cases.
- Attempting to use different slices on a disk and share that disk among different operating systems or with a different ZFS storage pool or storage pool components is not recommended.

## ▼ How to Recreate the ZFS Root Pool (EFI (GPT))

Use the following procedure if you need to recreate the ZFS root pool or if you want to create an alternate root pool. The `zpool create` command below automatically creates a EFI (GPT) labeled disk with the correct boot information.

### 1 Become an administrator.

### 2 Identify the disks for the root pool.

Use the `format` utility to identify the disks for the root pool.

```
format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
 0: c6t0d0 <SUN72G cyl 14087 alt 2 hd 24 sec 424>
 /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@0,0
 1: c6t1d0 <FUJITSU-MAV2073RCSUN72G-0301-68.37GB>
 /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@1,0
 2: c6t2d0 <FUJITSU-MAV2073RCSUN72G-0301-68.37GB>
 /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@2,0
 3: c6t3d0 <FUJITSU-MAV2073RCSUN72G-0301 cyl 14087 alt 2 hd 24 sec 424>
 /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/sd@3,0
Specify disk (enter its number):
```

### 3 Recreate the root pool.

```
zpool create -B rpool mirror c1t0d0 c2t0d0
```

If you want to create an alternate root pool, then using syntax similar to the following:

```
zpool create -B rpool2 mirror c1t0d0 c2t0d0
beadm create -p rpool2 solaris2
beadm activate -p rpool2 solaris2
```

### 4 Restore the root pool snapshots, if necessary.

For information about complete ZFS root pool recovery, see [Chapter 11, “Archiving Snapshots and Root Pool Recovery,”](#) in *Oracle Solaris 11.1 Administration: ZFS File Systems*.

## ▼ x86: How to Create a Disk Slice for a ZFS Root File System (VTOC)

In general, the root pool disk is installed automatically when the system is installed. If you need to replace a root pool disk or attach a new disk as a mirrored root pool disk, see the steps below.

For a full description of `fdisk` partitions, see [“x86: Guidelines for Creating an `fdisk` Partition” on page 229](#).

### 1 Become an administrator.

### 2 Offline and unconfigure the failed disk, if necessary.

Some hardware requires that you offline and unconfigure a disk before attempting the `zpool` replace operation to replace a failed disk. For example:

```
zpool offline rpool c8t1d0s0
cfgadm -c unconfigure c8::dsk/c8t1d0
```

### 3 Physically connect the new or replacement disk to the system, if necessary.

a. Physically remove the failed disk.

b. Physically insert the replacement disk.

c. Configure the replacement disk, if necessary. For example:

```
cfgadm -c configure c8::dsk/c28t1d0
```

On some hardware, you do not have to reconfigure the replacement disk after it is inserted.

### 4 Confirm that the disk is accessible by reviewing the format output.

For example, the `format` command shows 4 disks connected to this system.

```
format -e
AVAILABLE DISK SELECTIONS:
 1. c8t0d0 <Sun-STK RAID INT-V1.0 cyl 17830 alt 2 hd 255 sec 63>
 /pci@0,0/pci10de,375ef/pci108e,286e0/disk@0,0
 2. c8t1d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375ef/pci108e,286e0/disk@1,0
 3. c8t2d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375ef/pci108e,286e0/disk@2,0
 4. c8t3d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375ef/pci108e,286e0/disk@3,0
```

### 5 Select the disk to be used for the ZFS root pool.

```
Specify disk (enter its number): 1
selecting c8t1d0
[disk formatted]
.
.
format>
```

**6 Review the status of the fdisk partition.**

- If the disk has no fdisk partition, you will see a message similar to the following:

```
format> fdisk
No Solaris fdisk partition found.
```

If so, go to the next step to create an fdisk partition.

- If the disk has an EFI fdisk or some other partition type, go to the next step to create a Solaris fdisk partition.
- If the disk has a Solaris fdisk partition, go to step 9 to create a disk slice for the root pool.

**7 If necessary, create a Solaris fdisk partition by selecting the fdisk option.**

```
format> fdisk
No fdisk table exists. The default partition for the disk is:
```

```
 a 100% "SOLARIS System" partition
```

Type "y" to accept the default partition, otherwise type "n" to edit the partition table. **y**

**8 If the disk has an EFI fdisk partition, then you will need to create a Solaris fdisk partition.**

If you print the disk's partition table with the format utility, and you see the partition table refers to the first sector and the size, then this is an EFI partition. You will need to create a Solaris fdisk partition as follows:

**a. Select fdisk from the format options.**

```
format -e c8t1d0
selecting c8t1d0
[disk formatted]
format> fdisk
```

**b. Delete the existing EFI partition by selecting option 3, Delete a partition.**

```
Enter Selection: 3
Specify the partition number to delete (or enter 0 to exit): 1
Are you sure you want to delete partition 1? This will make all files and
programs in this partition inaccessible (type "y" or "n"). y
```

Partition 1 has been deleted.

**c. Create a new Solaris partition by selecting option 1, Create a partition.**

```
Enter Selection: 1
Select the partition type to create: 1
Specify the percentage of disk to use for this partition
(or type "c" to specify the size in cylinders). 100
Should this become the active partition? If yes, it will be activated
each time the computer is reset or turned on.
Please type "y" or "n". y
Partition 1 is now the active partition.
```

**d. Update the disk configuration and exit.**

```
Enter Selection: 6
format>
```

**e. Display the SMI partition table. If the default partition table is applied, then slice 0 might be 0 in size or it might be too small. See the next step.**

```
format> partition
partition> print
```

**9 Confirm that the disk has an SMI label by displaying the partition (slice) information and review the slice 0 size information.**

Set the free hog partition so that all the unallocated disk space is collected in slice 0. Then, press return through the slice size fields to create one large slice 0.

```
partition> modify
Select partitioning base:
 0. Current partition table (default)
 1. All Free Hog
Choose base (enter number) [0]? 1
Part Tag Flag Cylinders Size Blocks
 0 root wm 0 0 (0/0/0) 0
 1 swap wu 0 0 (0/0/0) 0
 2 backup wu 0 - 17829 136.58GB (17830/0/0) 286438950
 3 unassigned wm 0 0 (0/0/0) 0
 4 unassigned wm 0 0 (0/0/0) 0
 5 unassigned wm 0 0 (0/0/0) 0
 6 usr wm 0 0 (0/0/0) 0
 7 unassigned wm 0 0 (0/0/0) 0
 8 boot wu 0 - 0 7.84MB (1/0/0) 16065
 9 alternates wm 0 0 (0/0/0) 0
```

```
Do you wish to continue creating a new partition
table based on above table[yes]?
```

```
Free Hog partition[6]? 0
Enter size of partition '1' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '3' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '4' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '5' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '6' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '7' [0b, 0c, 0.00mb, 0.00gb]:
```

```
Part Tag Flag Cylinders Size Blocks
 0 root wm 1 - 17829 136.58GB (17829/0/0) 286422885
 1 swap wu 0 0 (0/0/0) 0
 2 backup wu 0 - 17829 136.58GB (17830/0/0) 286438950
 3 unassigned wm 0 0 (0/0/0) 0
 4 unassigned wm 0 0 (0/0/0) 0
 5 unassigned wm 0 0 (0/0/0) 0
 6 usr wm 0 0 (0/0/0) 0
 7 unassigned wm 0 0 (0/0/0) 0
 8 boot wu 0 - 0 7.84MB (1/0/0) 16065
 9 alternates wm 0 0 (0/0/0) 0
```

```
Do you wish to continue creating a new partition
table based on above table[yes]? yes
```

```
Enter table name (remember quotes): "c8t0d0"
```

```
Ready to label disk, continue? yes
```

**10 Let ZFS know that the failed disk is replaced.**

```
zpool replace rpool c8t1d0s0
zpool online rpool c8t1d0s0
```

On some hardware, you do not have to online the replacement disk after it is inserted.

If you are attaching a new disk to create a mirrored root pool or attaching a larger disk to replace a smaller disk, use syntax similar to the following:

```
zpool attach rpool c8t0d0s0 c8t1d0s0
```

A `zpool attach` operation on a root pool disk automatically applies the boot blocks.

**11 If a root pool disk is replaced with a new disk, apply the boot blocks.**

For example:

```
bootadm install-bootloader
```

A `zpool replace` operation does not automatically apply the boot blocks.

**12 Verify that you can boot from the new disk.****13 If the system boots from the new disk, detach the old disk.**

This step is only necessary if you attach a new disk to replace a failed disk or a smaller disk.

```
zpool detach rpool c8t0d0s0
```

**14 Set up the system to boot automatically from the new disk by reconfiguring the system's BIOS.****▼ x86: How to Replace a ZFS Root Pool Disk (EFI (GPT))**

In general, the root pool disk is installed automatically when the system is installed. If you need to replace a root pool disk or attach a new disk as a mirrored root pool disk, see the steps below.

In Oracle Solaris 11.1, in most cases, an EFI (GPT) disk label is installed on the root pool disk.

For a full description of `fdisk` partitions, see [“x86: Guidelines for Creating an `fdisk` Partition” on page 229](#).

**1 Become an administrator.****2 Offline and unconfigure the failed disk, if necessary.**

Some hardware requires that you offline and unconfigure a disk before attempting the `zpool replace` operation to replace a failed disk. For example:

```
zpool offline rpool c8t1d0
cfgadm -c unconfigure c8::disk/c8t1d0
```

**3 Physically connect the new or replacement disk to the system, if necessary.**

- a. Physically remove the failed disk.
- b. Physically insert the replacement disk.
- c. Configure the replacement disk, if necessary. For example:

```
cfgadm -c configure c8::disk/c8t1d0
```

On some hardware, you do not have to reconfigure the replacement disk after it is inserted.

**4 Confirm that the disk is accessible by reviewing the format output.**

For example, the format command sees 4 disks connected to this system.

```
format -e
AVAILABLE DISK SELECTIONS:
 1. c8t0d0 <Sun-STK RAID INT-V1.0 cyl 17830 alt 2 hd 255 sec 63>
 /pci@0,0/pci10de,375ef/pci108e,286@0/disk@0,0
 2. c8t1d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375ef/pci108e,286@0/disk@1,0
 3. c8t2d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375ef/pci108e,286@0/disk@2,0
 4. c8t3d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375ef/pci108e,286@0/disk@3,0
```

**5 Let ZFS know that the failed disk is replaced.**

```
zpool replace rpool c8t1d0
zpool online rpool c8t1d0
```

On some hardware, you do not have to online the replacement disk after it is inserted.

If you are attaching a new disk to create a mirrored root pool or attaching a larger disk to replace a smaller disk, use syntax similar to the following:

```
zpool attach rpool c8t0d0 c8t1d0
```

A `zpool attach` operation on a root pool disk applies the boot blocks automatically.

If your root pool disk contains customized partitions, you might need to use syntax similar to the following:

```
zpool attach rpool c8t0d0s0 c8t0d0
```

**6 If a root pool disk is replaced with a new disk, apply the boot blocks.**

For example:

```
bootadm install-bootloader
```

A `zpool replace` operation on a root pool disk does not apply the boot blocks automatically.

**7 Verify that you can boot from the new disk.**

**8 If the system boots from the new disk, detach the old disk.**

This step is only necessary if you attach a new disk to replace a failed disk or a smaller disk.

```
zpool detach rpool c8t0d0
```

**9 Set up the system to boot automatically from the new disk by reconfiguring the system's BIOS.****▼ x86: How to Replace a ZFS Root Pool Disk (VTOC)**

In general, the root pool disk is installed automatically when the system is installed. If you need to replace a root pool disk or attach a new disk as a mirrored root pool disk, see the steps below.

For a full description of `fdisk` partitions, see [“x86: Guidelines for Creating an `fdisk` Partition” on page 229](#).

**1 Become an administrator.****2 Offline and unconfigure the failed disk, if necessary.**

Some hardware requires that you offline and unconfigure a disk before attempting the `zpool replace` operation to replace a failed disk. For example:

```
zpool offline rpool c8t1d0
cfgadm -c unconfigure c8::dsk/c8t1d0
```

**3 Physically connect the new or replacement disk to the system, if necessary.**

a. Physically remove the failed disk.

b. Physically insert the replacement disk.

c. Configure the replacement disk, if necessary. For example:

```
cfgadm -c configure c8::dsk/c8t1d0
```

On some hardware, you do not have to reconfigure the replacement disk after it is inserted.

**4 Confirm that the disk is accessible by reviewing the format output.**

For example, the `format` command sees 4 disks connected to this system.

```
format -e
AVAILABLE DISK SELECTIONS:
 1. c8t0d0 <Sun-STK RAID INT-V1.0 cyl 17830 alt 2 hd 255 sec 63>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@0,0
 2. c8t1d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@1,0
 3. c8t2d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@2,0
 4. c8t3d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@3,0
```

**5 Select the disk to be used for the ZFS root pool.**

```
Specify disk (enter its number): 1
selecting c8t1d0
[disk formatted]
.
.
.
format>
```

**6 Review the status of the fdisk partition.**

- If the disk has no fdisk partition, you will see a message similar to the following:

```
format> fdisk
No Solaris fdisk partition found.
```

If so, go to step 4 to create an fdisk partition.

- If the disk has an EFI fdisk or some other partition type, go to the next step to create a Solaris fdisk partition.
- If the disk has a Solaris fdisk partition, go to step 9 to create a disk slice for the root pool.

**7 If necessary, create a Solaris fdisk partition by selecting the fdisk option.**

```
format> fdisk
No fdisk table exists. The default partition for the disk is:
```

```
 a 100% "SOLARIS System" partition
```

Type "y" to accept the default partition, otherwise type "n" to edit the partition table. **y**

**8 If the disk has an EFI fdisk partition, then you will need to create a Solaris fdisk partition.**

If you print the disk's partition table with the format utility, and you see the partition table refers to the first sector and the size, then this is an EFI partition. You will need to create a Solaris fdisk partition as follows:

- Select fdisk from the format options.

```
format -e c8t1d0
selecting c8t1d0
[disk formatted]
format> fdisk
```

- Delete the existing EFI partition by selecting option 3, Delete a partition.

```
Enter Selection: 3
Specify the partition number to delete (or enter 0 to exit): 1
Are you sure you want to delete partition 1? This will make all files and
programs in this partition inaccessible (type "y" or "n"). y
```

Partition 1 has been deleted.

- Create a new Solaris partition by selecting option 1, Create a partition.

```

Enter Selection: 1
Select the partition type to create: 1
Specify the percentage of disk to use for this partition
(or type "c" to specify the size in cylinders). 100
Should this become the active partition? If yes, it will be activated
each time the computer is reset or turned on.
Please type "y" or "n". y
Partition 1 is now the active partition.

```

- Update the disk configuration and exit.

```

Enter Selection: 6
format>

```

- Display the SMI partition table. If the default partition table is applied, then slice 0 might be 0 in size or it might be too small. See the next step.

```

format> partition
partition> print

```

## 9 Confirm that the disk has an SMI label by displaying the partition (slice) information and review the slice 0 size information.

Set the free hog partition so that all the unallocated disk space is collected in slice 0. Then, press return through the slice size fields to create one large slice 0.

```

partition> modify
Select partitioning base:
 0. Current partition table (default)
 1. All Free Hog
Choose base (enter number) [0]? 1

```

Part	Tag	Flag	Cylinders	Size	Blocks
0	root	wm	0	0	(0/0/0) 0
1	swap	wu	0	0	(0/0/0) 0
2	backup	wu	0 - 17829	136.58GB	(17830/0/0) 286438950
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0
5	unassigned	wm	0	0	(0/0/0) 0
6	usr	wm	0	0	(0/0/0) 0
7	unassigned	wm	0	0	(0/0/0) 0
8	boot	wu	0 - 0	7.84MB	(1/0/0) 16065
9	alternates	wm	0	0	(0/0/0) 0

```

Do you wish to continue creating a new partition
table based on above table[yes]?

```

```

Free Hog partition[6]? 0
Enter size of partition '1' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '3' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '4' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '5' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '6' [0b, 0c, 0.00mb, 0.00gb]:
Enter size of partition '7' [0b, 0c, 0.00mb, 0.00gb]:

```

Part	Tag	Flag	Cylinders	Size	Blocks
0	root	wm	1 - 17829	136.58GB	(17829/0/0) 286422885
1	swap	wu	0	0	(0/0/0) 0
2	backup	wu	0 - 17829	136.58GB	(17830/0/0) 286438950
3	unassigned	wm	0	0	(0/0/0) 0
4	unassigned	wm	0	0	(0/0/0) 0

```

5 unassigned wm 0 0 (0/0/0) 0
6 usr wm 0 0 (0/0/0) 0
7 unassigned wm 0 0 (0/0/0) 0
8 boot wu 0 - 0 7.84MB (1/0/0) 16065
9 alternates wm 0 0 (0/0/0) 0

```

Do you wish to continue creating a new partition

table based on above table[yes]? **yes**

Enter table name (remember quotes): **"c8t1d0"**

Ready to label disk, continue? **yes**

## 10 Let ZFS know that the failed disk is replaced.

```
zpool replace rpool c8t1d0s0
```

```
zpool online rpool c8t1d0s0
```

On some hardware, you do not have to online the replacement disk after it is inserted.

If you are attaching a new disk to create a mirrored root pool or attaching a larger disk to replace a smaller disk, use syntax similar to the following:

```
zpool attach rpool c8t0d0s0 c8t1d0s0
```

When using the `zpool attach` command on a root pool, the boot blocks are applied automatically.

## 11 If a root pool disk is replaced with a new disk, apply the boot blocks.

For example:

```
bootadm install-bootloader
```

## 12 Verify that you can boot from the new disk.

## 13 If the system boots from the new disk, detach the old disk.

This step is only necessary if you attach a new disk to replace a failed disk or a smaller disk.

```
zpool detach rpool c8t1d0s0
```

## 14 Set up the system to boot automatically from the new disk by reconfiguring the system's BIOS.

## ▼ x86: How to Install Boot Blocks for a ZFS Root File System

If you replace a root pool disk with the `zpool replace` command, you must install the boot loader. The following procedures works for both VTOC and EFI (GPT) labels.

### 1 Become an administrator.

### 2 Install the boot blocks on the system disk.

```
bootadm install-bootloader
```

If you need to install the boot loader on an alternate root pool, then use the `-P (pool)` option.

```
bootadm install-bootloader -P rpool2
```

If you want to install the GRUB Legacy boot loader, you must first remove all GRUB 2 boot environments from your system and then use the `installgrub` command. For instructions, see [“Installing GRUB Legacy on a System That Has GRUB 2 Installed”](#) in *Booting and Shutting Down Oracle Solaris 11.1 Systems*.

### 3 Verify that the boot blocks are installed by rebooting the system to run level 3.

```
init 6
```

## ▼ x86: How to Set Up a Disk for a ZFS Non-Root File System

If you are setting up a disk to be used with a non-root ZFS file system, the disk is relabeled automatically when the pool is created or when the disk is added to the pool. If a pool is created with whole disks or when a whole disk is added to a ZFS storage pool, an EFI label is applied. For more information about EFI disk labels, see [“EFI \(GPT\) Disk Label”](#) on page 163.

Generally, most modern bus types support hot-plugging. This means you can insert a disk in an empty slot and the system recognizes it. For more information about hot-plugging devices, see [Chapter 4, “Dynamically Configuring Devices \(Tasks\)”](#).

### 1 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights”](#) in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Connect the disk to the system and check the disk's physical connections.

Refer to the disk's hardware installation guide for details.

### 3 Offline and unconfigure the failed disk, if necessary.

Some hardware requires that you offline and unconfigure a disk before attempting the `zpool replace` operation to replace a failed disk. For example:

```
zpool offline tank c1t1d0
cfgadm -c unconfigure c1::dsk/c1t1d0
<Physically remove failed disk c1t1d0>
<Physically insert replacement disk c1t1d0>
cfgadm -c configure c1::dsk/c1t1d0
```

On some hardware, you do not to reconfigure the replacement disk after it is inserted.

### 4 Confirm that the new disk is recognized.

Review the output of the `format` utility to see if the disk is listed under `AVAILABLE DISK SELECTIONS`. Then, quit the `format` utility.

```
format
```

**5 Let ZFS know that the failed disk is replaced, if necessary.**

```
zpool replace tank c1t1d0
zpool online tank c1t1d0
```

Confirm that the new disk is resilvering.

```
zpool status tank
```

**6 Attach a new disk to an existing ZFS storage pool, if necessary.**

For example:

```
zpool attach tank mirror c1t0d0 c2t0d0
```

Confirm that the new disk is resilvering.

```
zpool status tank
```

For more information, see [Chapter 3, “Managing Oracle Solaris ZFS Storage Pools,”](#) in *Oracle Solaris 11.1 Administration: ZFS File Systems*.

## x86: Creating and Changing Solaris fdisk Partitions

Review the following sections for guidelines and examples of creating or changing Solaris fdisk partitions.

### x86: Guidelines for Creating an fdisk Partition

Follow these guidelines when you set up one or more fdisk partitions:

- The disk can be divided into a maximum of four fdisk partitions. One of partitions must be a Solaris partition.
- The Solaris partition must be made the active partition on the disk. The active partition is partition whose operating system will be booted by default at system startup.
- Solaris fdisk partitions must begin on cylinder boundaries.
- Solaris fdisk partitions must begin at cylinder 1, not cylinder 0, on the first disk because additional boot information, including the master boot record, is written in sector 0.
- The Solaris fdisk partition can be the entire disk. Or, you might want to make it smaller to allow room for a DOS partition. You can also make a new fdisk partition on a disk without disturbing existing partitions (if sufficient space is available) to create a new partition.

**x86 only** – Solaris slices are also called partitions. Certain interfaces might refer to a *slice* as a *partition*.

fdisk partitions are supported only on x86 based systems. To avoid confusion, Oracle Solaris documentation tries to distinguish between fdisk partitions and the entities within the Solaris fdisk partition. These entities might be called slices or partitions.

---

## ▼ x86: How to Create a Solaris fdisk Partition

**Before You Begin** If you need information about fdisk partitions, see “[x86: Guidelines for Creating an fdisk Partition](#)” on page 229.

### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Invoke the format utility.

```
format
```

A numbered list of disks is displayed.

For more information, see [format\(1M\)](#).

### 3 Type the number of the disk on which to create a Solaris fdisk partition.

Specify disk (enter its number): *disk-number*

where *disk-number* is the number of the disk on which you want to create a Solaris fdisk partition.

### 4 Select the fdisk menu.

```
format> fdisk
```

The fdisk menu that is displayed depends upon whether the disk has existing fdisk partitions.

### 5 Create and activate a Solaris fdisk partition that spans the entire disk by specifying y at the prompt. Then, go to step 13.

No fdisk table exists. The default partition for the disk is:

```
a 100% "SOLARIS System" partition
```

Type "y" to accept the default partition, otherwise type "n" to edit the partition table.

```
y
```

**6 Specify n at the prompt if you do not want the Solaris fdisk partition to span the entire disk.**

Type "y" to accept the default partition, otherwise type "n" to edit the partition table.

```
n
 Total disk size is 17848 cylinders
 Cylinder size is 16065 (512 byte) blocks

 Cylinders
Partition Status Type Start End Length %
===== ===== =====

```

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Edit/View extended partitions
6. Exit (update disk configuration and exit)
7. Cancel (exit without updating disk configuration)

Enter Selection:

**7 Select option 1, Create a partition, to create an fdisk partition.**

Enter Selection: 1

**8 Create a Solaris fdisk partition by selecting 1 (=Solaris2).**

Select the partition type to create:

```
1=SOLARIS2 2=UNIX 3=PCIX05 4=Other 5=DOS12
6=DOS16 7=DOSEXT 8=DOSBIG 9=DOS16LBA A=x86 Boot
B=Diagnostic C=FAT32 D=FAT32LBA E=DOSEXTLBA F=EFI (Protective)
G=EFI_SYS 0=Exit?
```

**9 Identify the percentage of the disk to be reserved for the Solaris fdisk partition. Keep in mind the size of any existing fdisk partitions when you calculate this percentage.**

Specify the percentage of disk to use for this partition (or type "c" to specify the size in cylinders). *nn*

**10 Activate the Solaris fdisk partition by typing y at the prompt.**

Should this become the active partition? If yes, it will be activated each time the computer is reset or turned on. Please type "y" or "n". **y**

The Enter Selection prompt is displayed after the fdisk partition is activated.

**11 Select option 1, Create a partition, to create another fdisk partition.**

See steps 8–10 for instructions on creating an fdisk partition.

**12 Update the disk configuration, and exit the fdisk menu from the selection menu.**

Selection: 6

**13 Relabel the disk by using the `label` command.**

```
format> label
Ready to label disk, continue? yes
format>
```

**14 Quit the format utility.**

```
format> quit
```

**Example 10–2 x86: Creating a Solaris fdisk Partition That Spans the Entire Drive**

The following example uses the `format` utility's `fdisk` option to create a Solaris fdisk partition that spans the entire drive.

```
format
Searching for disks...done
AVAILABLE DISK SELECTIONS:
 0. c8t0d0 <SEAGATE-ST973401LSUN72G-0556 cyl 8921 alt 2 hd 255 sec 63>
 /pci@0,0/pci1022,7458@11/pci1000,3060@4/sd@0,0
 /dev/chassis/SYS/HD0/disk
 1. c8t1d0 <SEAGATE-ST973401LSUN72G-0556 cyl 8921 alt 2 hd 255 sec 63>
 /pci@0,0/pci1022,7458@11/pci1000,3060@4/sd@1,0
 /dev/chassis/SYS/HD1/disk
 2. c8t2d0 <SEAGATE-ST973401LSUN72G-0556-68.37GB>
 /pci@0,0/pci1022,7458@11/pci1000,3060@4/sd@2,0
 /dev/chassis/SYS/HD2/disk
 3. c8t3d0 <SEAGATE-ST973401LSUN72G-0556 cyl 14087 alt 2 hd 24 sec 424>
 /pci@0,0/pci1022,7458@11/pci1000,3060@4/sd@3,0
 /dev/chassis/SYS/HD3/disk

Specify disk (enter its number): 3
selecting c8t3d0
[disk formatted]
No Solaris fdisk partition found.
format> fdisk
No fdisk table exists. The default partitioning for your disk is:

 a 100% "SOLARIS System" partition.

Type "y" to accept the default partition, otherwise type "n" to edit the
partition table. y

format> label
Ready to label disk, continue? yes
format> quit
```

## Changing the fdisk Partition Identifier

The Solaris fdisk partition identifier on x86 systems has been changed from 130 (0x82) to 191 (0xbf). All Oracle Solaris commands, utilities, and drivers have been updated to work with either fdisk identifier. There is no change in fdisk functionality.

## ▼ How to Change the Solaris fdisk Identifier

A new fdisk menu option enables you to switch back and forth between the new and old identifier. The fdisk identifier can be changed even when the file system that is contained in the partition is mounted.

Two type values in the fdisk menu reflect the old and new identifiers as follows:

- Solaris identifies 0x82
- Solaris2 identifies 0xbf

- 1 **Become an administrator.**
- 2 **Display the current fdisk identifier.**

For example:

```
Total disk size is 39890 cylinders
Cylinder size is 4032 (512 byte) blocks
```

Partition	Status	Type	Cylinders		Length	%
			Start	End		
1	Active	x86 Boot	1	6	6	0
2		Solaris2	7	39889	39883	100

- 3 **Select option 4 from the fdisk menu to change the fdisk partition identifier back to 0x82.**

SELECT ONE OF THE FOLLOWING:

1. Create a partition
2. Specify the active partition
3. Delete a partition
4. Change between Solaris and Solaris2 Partition IDs
5. Exit (update disk configuration and exit)
6. Cancel (exit without updating disk configuration)

Enter Selection: 4

- 4 **Select option 6 to update your disk configuration and exit.**
- 5 **If necessary, select option 4 from the fdisk menu to change the fdisk partition identifier back to 0xbf.**

For example:

```
Total disk size is 39890 cylinders
Cylinder size is 4032 (512 byte) blocks
```

Partition	Status	Type	Cylinders		Length	%
			Start	End		
1	Active	x86 Boot	1	6	6	0
2		Solaris	7	39889	39883	100

SELECT ONE OF THE FOLLOWING:

1. Create a partition

2. Specify the active partition
  3. Delete a partition
  4. Change between Solaris and Solaris2 Partition IDs
  5. Exit (update disk configuration and exit)
  6. Cancel (exit without updating disk configuration)
- Enter Selection: **4**

**6 Select option 6 to update your disk configuration and exit.**

# Configuring Storage Devices With COMSTAR (Tasks)

---

This chapter describes how to configure Common Multiprotocol SCSI TARget, or COMSTAR, a software framework that enables you to convert any Oracle Solaris 11 host into a SCSI target device that can be accessed over a storage network by initiator hosts.

This means you can make storage devices on a system available to Linux, Mac OS, or Windows client systems as if they were local storage devices. Supported storage protocols are iSCSI, FC, iSER, and SRP.

For information about configuring iSCSI initiators, see [Chapter 4, “Configuring Solaris iSCSI Initiators,”](#) in *Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing*.

For information about the iSNS support in Oracle Solaris, see [Chapter 12, “Configuring and Managing the Oracle Solaris Internet Storage Name Service \(iSNS\).”](#)

For troubleshooting general iSCSI configuration problems in Oracle Solaris, see [“Troubleshooting iSCSI Configuration Problems”](#) in *Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing*.

## Overview of COMSTAR Features

COMSTAR utilizes a SCSI Target Mode Framework (STMF) to manage target storage devices with the following components:

- Port providers (or plug-ins) – Implement protocols, such as Fibre Channel (FC) and iSCSI.
- Logical unit providers – Emulate various SCSI devices, such as disk and tape devices.
- The `libstmf` management library – Provides the COMSTAR management interface. The modules that implement the iSCSI functionality do not interact directly with the underlying transport. In a similar way, the modules that implement the transport protocol are unaware of the SCSI-level functionality that is inherent in the packets they are transporting. Some

transport examples are Fibre Channel and iSCSI. The framework separates the execution and cleanup of SCSI commands and the associated resources. This separation simplifies the task of writing SCSI or transport modules.

Use the following to administer these features:

- The `itadm` command manages Internet SCSI (iSCSI) nodes within the SCSI target mode framework.
- The `stmfadm` command configures logical units within the SCSI target mode framework.
- The `srptadm` command manages SCSI RDMA Protocol (SRP) target ports within the SCSI target mode framework.

## Identifying COMSTAR Software and Hardware Requirements

- Oracle Solaris storage software and devices
- The `group/feature/storage-server` software package for the system that provides the storage devices
- Any supported NIC

## Configuring Storage Devices With COMSTAR (Task Map)

This is a general list of tasks associated with configuring storage devices with COMSTAR. Some of the tasks are optional depending on your network configuration needs. Some of the links below will take you to separate documents that describe network configuration and initiator configuration.

- [“Identifying COMSTAR Software and Hardware Requirements” on page 236](#)
- [Introduction to Oracle Solaris 11 Networking](#)
- [“Configuring Storage Devices With COMSTAR \(Tasks\)” on page 238](#)
- [“Configuring Dynamic or Static Target Discovery” on page 237](#)
- Chapter 4, “Configuring Solaris iSCSI Initiators,” in *Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing*
- [“How to Access iSCSI Disks” on page 245](#)
- [“Configuring Authentication in Your iSCSI-Based Storage Network” in Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing](#)
- [“Setting Up iSCSI Multipathed Devices in Oracle Solaris” in Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing](#)
- [“Monitoring Your iSCSI Configuration” in Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing](#)

## COMSTAR Terminology

Review the following terminology before configuring target devices with COMSTAR.

Term	Description
Discovery	The process that presents the initiator with a list of available targets.
Discovery method	The way in which the iSCSI targets can be found. Three methods are currently available: <ul style="list-style-type: none"> <li>■ Internet Storage Name Service (iSNS) – Potential targets are discovered by interacting with one or more iSNS servers.</li> <li>■ SendTargets – Potential targets are discovered by using a <i>discovery-address</i>.</li> <li>■ Static – Static target addressing is configured.</li> </ul>
Initiator	The driver that initiates SCSI requests to the iSCSI target.
Initiator group	A set of initiators. When an initiator group is associated with a , only initiators from that group may access the .
iqn or eui address format	An iqn (iSCSI qualified name) address is the unique identifier for a device in an iSCSI network using the form <i>iqn.date.authority:uniqueid</i> . An iSCSI initiator or target is assigned an IQN name automatically when the iSCSI initiator or target is initialized.  An eui (extended unique identifier) address consists of 16 hexadecimal digits, and identifies a class of GUIDs that is used in both the SCSI and InfiniBand standards. SRP devices use the eui address format.
Logical unit	A uniquely numbered component in a storage system. When a is associated with one or more SCSI targets, the target can be accessed by one or more SCSI initiators.
Target device	The iSCSI storage component.
Target group	A set of targets. A can be made available to all targets in one target group.
Target portal group	A list of IP addresses that determines which interfaces a specific iSCSI target will listen to. A TPG contains IP addresses and TCP port numbers

## Configuring Dynamic or Static Target Discovery

Determine whether you want to configure one of the dynamic device discovery methods or use static iSCSI initiator targets to perform device discovery.

- **Dynamic device discovery** – Two dynamic device discovery methods are available:

- **SendTargets** – If an iSCSI node exposes a large number of targets, such as an iSCSI to Fibre-Channel bridge, you can supply the iSCSI node IP address/port combination and allow the iSCSI initiator to use the SendTargets features to perform the device discovery.
- **iSNS** – The Internet Storage Name Service (iSNS) allows the iSCSI initiator to discover the targets to which it has access using as little configuration information as possible. It also provides state change notification to notify the iSCSI initiator when changes in the operational state of storage nodes occur. To use the iSNS discovery method, you can supply the iSNS server address/port combination and allow the iSCSI initiator to query the iSNS servers that you specified to perform the device discovery. The default port for the iSNS server is 3205. For more information about iSNS, see RFC 4171:

<http://www.ietf.org/rfc/rfc4171.txt>

The iSNS discovery service provides an administrative model to discover all targets on a network.

For more information about setting up iSNS support in Oracle Solaris, see [Chapter 12, “Configuring and Managing the Oracle Solaris Internet Storage Name Service \(iSNS\).”](#)

- **Static device discovery** – If an iSCSI node has few targets or if you want to restrict the targets that the initiator attempts to access, you can statically configure the *target-name* by using the following static target address naming convention:

*target,target-address[:port-number]*

You can determine the static target address from the array's management tool.

---

**Note** – Do not configure an iSCSI target to be discovered by both static and dynamic device discovery methods. The consequence of using redundant discovery methods might be slow performance when the initiator is communicating with the iSCSI target device.

---

## Configuring Storage Devices With COMSTAR (Tasks)

You can set up and configure a COMSTAR Internet SCSI (iSCSI) target and make it available over the network. The iSCSI features can work over a normal Internet connection (such as Ethernet) using the standard iSCSI protocol. The iSCSI protocol also provides naming and discovery services, authentication services using CHAP and RADIUS, and centralized management through iSNS.

If the system has at least one InfiniBand (IB) Host Channel Adapter (HCA) and connectivity exists between the initiator and the target, the iSCSI connection uses iSCSI Extensions for RDMA (iSER) for enhanced data throughput. iSER used with an IB HCA provides high bandwidth, low CPU utilization, and a single network connection that multiple protocols can share.

The iSER target and initiator both use a component called iSCSI Data Mover to provide iSCSI connection services. Use of iSER is automatic whenever both the initiator and the target are configured to use IP addresses that correspond to IB-capable devices.

## ▼ How to Enable the STMF Service

COMSTAR uses SMF to store its current, persistent configuration, such as logical unit mapping, host group definitions, and target group definitions. When the service is enabled during boot or when using the `svcadm` command, it clears any stale configuration data inside the kernel framework, and then reloads the configuration from the SMF repository into the driver. After the configuration is loaded, any changes that are made to the configuration are automatically updated inside the driver database, as well as inside the SMF repository. For example, any changes made through the `stmfadm` command are automatically updated in both areas.

The COMSTAR target mode framework runs as the `stmf` service. By default, the service is disabled. You must enable the service to use COMSTAR functionality. You can identify the service with the `svcs` command. If you have not rebooted the server since installing the `group/feature/storage-server` package, the service might not be enabled correctly.

### 1 Install the COMSTAR storage server software.

```
target# pkg install group/feature/storage-server
 Packages to install: 75
 Create boot environment: No
 Services to restart: 7
DOWNLOAD PKGS FILES XFER (MB)
Completed 75/75 9555/9555 105.7/105.7

PHASE ACTIONS
Install Phase 13347/13347

PHASE ITEMS
Package State Update Phase 75/75
Image State Update Phase 2/2
Loading smf(5) service descriptions: 17/17
Loading smf(5) service descriptions: 3/3

PHASE ITEMS
Reading Existing Index 8/8
Indexing Packages 75/75
Indexing Packages 75/75
Optimizing Index...

PHASE ITEMS
Indexing Packages 573/573
```

### 2 Either reboot the system or enable the `stmf` service.

```
target# svcadm enable stmf
svcs stmf
STATE STIME FMRI
online 09:42:32 svc:/system/stmf:default
```

## ▼ How to Back Up and Restore a COMSTAR Configuration

After you complete your COMSTAR configuration, make a copy that can be restored, if needed.

- 1 **Become an administrator.**
- 2 **Export the current COMSTAR configuration.**  

```
svccfg export -a stmf > COMSTAR.backup
```
- 3 **If necessary, restore the exported configuration.**  

```
svccfg import COMSTAR.backup
```

## ▼ How to Create a Logical Unit

The logical unit provider for creating disk-type LUs is called `sbd`. However, you must initialize the storage for the logical unit before you can share a disk-type LU.

The disk volume provided by the server is referred to as the *target*. When the LU is associated with an iSCSI target, it can be accessed by an iSCSI initiator.

The process for creating SCSI LUs is as follows:

- Initialize the storage for the LU, also known as the *backing store*.
- Create a SCSI LU by using the backing store.

When a LU is created, it is assigned a global unique identifier (GUID), for example, `600144F0B5418B0000004DDAC7C10001`. The GUID is used to refer to the LU in subsequent tasks, such as mapping a LU to select hosts.

The following steps are completed on the system that is providing the storage device.

- 1 **Create a ZFS storage pool.**  

```
target# zpool create sanpool mirror c2t3d0 c2t4d0
```
- 2 **Create a ZFS volume to be used as a SCSI LU.**  

```
target# zfs create -V 2g sanpool/vol1
```
- 3 **Create a LU for the ZFS volume.**  

```
target# stmfadm create-lu /dev/zvol/rdisk/sanpool/vol1
Logical unit created: 600144F0B5418B0000004DDAC7C10001
```

You can find the device path for the ZFS volume in the `/dev/zvol/rdisk/pool-name/` directory.
- 4 **Confirm that the LU has been created.**  

```
target# stmfadm list-lu
LU Name: 600144F0B5418B0000004DDAC7C10001
```

## 5 Add the LU view.

This command makes the LU accessible to all systems.

```
target# stmfadm add-view 600144F0B5418B0000004DDAC7C10001
```

If you want to restrict the LU view to specific systems, see [“How to Restrict Logical Unit Access to Selected Systems”](#) on page 247.

## 6 Verify the LU configuration.

```
target# stmfadm list-view -l 600144F0B5418B0000004DDAC7C10001
View Entry: 0
 Host group : All
 Target group : All
 LUN : 0
```

# ▼ How to Create an iSCSI Target

This procedure assumes that you are logged in to the local system will contains the iSCSI target.

## 1 Enable the iSCSI target service.

```
target# svcadm enable -r svc:/network/iscsi/target:default
```

Confirm that the service is enabled.

```
target# svcs -l iscsi/target
fmri svc:/network/iscsi/target:default
name iscsi target
enabled true
state online
next_state none
state_time Mon May 23 14:48:59 2011
logfile /var/svc/log/network-iscsi-target:default.log
restarter svc:/system/svc/restarter:default
dependency require_any/error svc:/milestone/network (online)
dependency require_all/none svc:/system/stmf:default (online)
```

## 2 Create the iSCSI target.

```
target# itadm create-target
Target iqn.1986-03.com.sun:02:73d12edc-9bb9-cb44-efc4-c3b36c039405 successfully created
```

## 3 Display the iSCSI target information.

```
target# itadm list-target -v
TARGET NAME STATE SESSIONS
iqn.1986-03.com.sun:02:73d12edc-9bb9-cb44-efc4-c3b36c039405 online 0
 alias: -
 auth: none (defaults)
 targetchapuser: -
 targetchapsecret: unset
 tpg-tags: default
```

## ▼ How to Enable iSNS Discovery for the Target Device

Use this method if you have at least one iSNS server on your network. This method enables the iSCSI initiator to discover targets with as little configuration as possible, while also providing state-change notification. State-change functionality notifies the iSCSI initiator when there are changes in the operational state of a target. For this method, you must supply the iSNS server address and port combination. The iSCSI initiator can then query the specified iSNS server to perform device discovery. The default port for an iSNS server is 3205.

### 1 Become an administrator.

### 2 Add the iSNS server information.

```
target# itadm modify-defaults --isns-server ip-address
```

Identify the *ip-address* of the iSNS server in your network.

This step adds the iSNS server information to all of the Oracle Solaris iSCSI targets.

### 3 Enable iSNS server discovery.

```
target# itadm modify-defaults --isns enable
```

## ▼ How to Configure an IB HCA for iSER

An InfiniBand (IB) Host Channel Adapter (HCA) is required to take full advantage of the iSCSI Extensions for RDMA (iSER) capabilities. To use iSER, you must configure the HCA on both the target and the initiator.

### 1 Become an administrator on the host (initiator) system.

### 2 Connect the HCA to an IB switch.

See the vendor documentation for details.

### 3 Check the InfiniBand port status.

```
dladm show-ib
```

LINK	HCAGUID	PORTGUID	PORT	STATE	PKEYS
net4	211700013E84C2	211700013E84C3	1	up	FFFF
net5	211700013E84C2	211700013E84C4	2	up	FFFF

### 4 Create the IB partition link.

```
dladm create-part -l net4 -P ffff ibd0
```

After the partition link is created, display the partition info.

```
dladm show-part
LINK PKEY OVER STATE FLAGS
ibd0 FFFF net4 unknown ----
```

##### 5 Configure the target and the initiator for the HCA.

The target and the initiator must be on the same subnet. This example uses `ibd0` as the driver.

```
ipadm create-ip ibd0
```

##### 6 Configure the IP address and port combination for the HCA.

```
ipadm create-addr -T static -a 100.100.10.100/24 ibd0/static
```

##### 7 Verify the interface configuration.

```
ipadm show-addr
```

##### 8 Become an administrator on the target system and repeat above steps for all other HCA hosts on the network.

##### 9 Verify connectivity on both the target and the initiator.

```
target# ping initiator-ip
initiator# ping target-ip
```

## Creating iSCSI Target Portal Groups

You can create a target portal group (TPG) to manage the discovery of multiple iSCSI and iSER targets. A TPG is a list of IP addresses to determine upon which interfaces a specific iSCSI target will listen.

A TPG contains IP addresses and TCP port numbers. To use this capability, you need to do the following:

- Create a TPG as a list of *ip-address:port* specifiers by using the `itadm create-tpg` command.
- Bind a specific iSCSI target to a TPG by using the `itadm modify-target -t` command.
- When an iSCSI target is made active, an iSCSI listener is created for each IP address and port belonging to a TPG associated with that target.

A TPG is an efficient way to control which targets are discovered through specific ports. For example, you could restrict your iSCSI target so that it is available only through one specific IP address or only through a set of iSER-capable IP addresses.

---

**Note** – Do not confuse target portal groups with target groups. A target group is a list of SCSI target ports that are all treated the same when creating views. Creating a view can help you facilitate LU mapping. Each view entry specifies a target group, a host group, and a LU. For more information on Target Groups and LUN mapping, see [“Making SCSI Logical Units Available” on page 245](#) and `stmfadm(1M)`.

---

To learn about static and iSNS target discovery, see [“Configuring Dynamic or Static Target Discovery” on page 237](#). The iSCSI initiator uses the `iscsiadm` command to discover TPGs. For more information, see `iscsiadm(1M)` and `itadm(1M)`.

## Using TPGs with iSER

When you use the SendTargets discovery and iSER at the same time, a common convention is to use a TPG to associate a specific iSCSI target port with only iSER-capable IP addresses. For example, if a target system has four IP addresses, A, B, C, and D, and only addresses B and C are iSER-capable, then addresses B and C could be added to a TPG, and assigned to a target T.

An iSCSI initiator with both Ethernet and InfiniBand (IB) interfaces could use the SendTargets discovery method to discover the possible storage targets. Without the use of TPGs, the initiator might always prefer the use of the Ethernet interfaces over the IB interfaces. By associating target T only with the IB interfaces, the initiator correctly prefers using its IB-capable interface when connecting to target T.

### ▼ How to Create a Target Portal Group for iSCSI Targets

You can create a target portal group (TPG) by providing a unique name, and a TPG Tag (ranging from 2–65535) is automatically generated. TPG Tag 1 is reserved for the default TPG that is used when you do not explicitly set a TPG on the target. The portal for the default TPG matches requests from all network interfaces on port 3260.

The following steps shows how to create two TPGs, TPGA and TPGB, that use port 8000 for the IP addresses in TPGB.

- 1 **Become an administrator.**
- 2 **Create two TPGs.**

```
target# itadm create-tpg TPGA 192.168.0.1 192.168.0.2
target# itadm create-tpg TPGB 192.168.0.2:8000 192.168.0.2:8000
```

---

**Note** – IPv4 portals are specified in dotted address notation (for example, 192.168.0.1). IPv6 portal addresses must be enclosed in square brackets.

---

- 3 **Configure an existing iSCSI target to use the TPGs, TPGA and TPGB.**

```
itadm modify-target -t TPGA,TPGB eui.20387ab8943ef7548
```

#### 4 Verify the TPGs that you created.

```
itadm list-tpg -v
```

You can remove a TPG with the `itadm delete-tpg` command.

## ▼ How to Access iSCSI Disks

After the devices have been discovered by the Oracle Solaris iSCSI initiator, the login negotiation occurs automatically. The Oracle Solaris iSCSI driver determines the number of available LUs and creates the device nodes. Then, the iSCSI devices can be treated as any other SCSI device.

You can create a ZFS storage pool on the LU and then create a ZFS file system.

You can view the iSCSI disks on the local system by using the `format` utility.

#### 1 Review the iSCSI LU information in the `format` output.

```
initiator# format
0. c0t600144F0B5418B0000004DDAC7C10001d0 <SUN-COMSTAR-1.0 cyl 1022 alt 2 hd 128 sec 32>
 /scsi_vhci/disk@g600144f0b5418b0000004ddac7c10001
1. c8t0d0 <Sun-STK RAID INT-V1.0 cyl 17830 alt 2 hd 255 sec 63>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@0,0
2. c8t1d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@1,0
3. c8t2d0 <Sun-STK RAID INT-V1.0-136.61GB>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@2,0
4. c8t3d0 <Sun-STK RAID INT-V1.0 cyl 17830 alt 2 hd 255 sec 63>
 /pci@0,0/pci10de,375@f/pci108e,286@0/disk@3,0
Specify disk (enter its number): 0
selecting c0t600144F0B5418B0000004DDAC7C10001d0
[disk formatted]
```

In the above output, disk 0 is an iSCSI LU under MPxIO control. Disks 1-4 are local disks.

#### 2 You can create a ZFS storage pool and ZFS file systems on the iSCSI LU.

```
initiator# zpool create pool-name c0t600144F0B5418B0000004DDAC7C10001d0
initiator# zfs create pool-name/fs-name
```

The ZFS file system is automatically mounted when created and is remounted at boot time.

## Making SCSI Logical Units Available

- [“How to Make a Logical Unit Available to All Systems” on page 246](#)
- [“How to Restrict Logical Unit Access to Selected Systems” on page 247](#)

Simply registering a logical unit (LU) with the STMF framework does not make it available to hosts (initiators) on the network. This section describes how to make LUs visible to initiator hosts for the following configurations.

For iSCSI, Fibre Channel, and FCoE configurations, a LU must be mapped before it can be accessed. You can choose one of the following methods, both of which use the `stmfadm` command:

- **Simple mapping** – Exposes the LU to all initiators through all the ports, using one command. Making LUs available to all hosts uses this method.
- **Selective mapping**– Enables you to specify the hosts that can access the LU. Making LUs available to selected hosts uses this method. This process includes the following steps:
  1. Defining host groups – A host group is a name given to a set of hosts (initiators) that are allowed to access the same LUs. This step is not needed if the same set of LUs is visible to all the hosts, as in simple mapping.
  2. Defining target groups – A target group is a name given to a set of target ports that export the same set of LUs to the same set of host groups. This step is not needed if the same set of LUs is visible to all the hosts, as in simple mapping.
  3. Adding one or more views for each logical unit – Adding a view creates a bridge between the LU and the host initiator. When an initiator from the host group logs in to a target port contained in the target group, the LU is visible.

---

**Note** – Don't confuse a target group with a target portal group (TPG). A TPG is a list of IP addresses that an iSCSI target listens to. A TPG can help you restrict an iSCSI target so that it is available only through one specific IP address. For more information on target groups, see [stmfadm\(1M\)](#).

---

A view entry consists of four components: host group, target group, logical unit number (LUN), and LU identifier. Of these four components, only the LU identifier is required. If the other components are omitted, the following default values are assumed:

- If the host group is omitted, the all initiators value is assumed.
- If the target group is omitted, the all targets value is assumed.
- If the LUN is omitted, the system chooses a suitable LUN for the entry.

## ▼ How to Make a Logical Unit Available to All Systems

This procedure makes a LU available to all initiator hosts on a storage network.

### 1 Obtain the Global Unique Identification (GUID) number for the LU.

```
stmfadm list-lu -v
```

### 2 Add a view for the logical unit.

```
3 # stmfadm add-view GUID-number
```

## ▼ How to Restrict Logical Unit Access to Selected Systems

Use this procedure to restrict LU access to selected hosts on a storage network. If you are using Fibre Channel ports, first identify the World Wide Names (WWN). Then, selectively map a logical unit number (LUN) to the ports on host -a, for example. A target group (targets-0) is also defined for a given set of target ports that export the same set of LUs to the same host group.

For information about configuring a target group, see [stmfadm\(1M\)](#).

- 1 **Become an administrator.**
- 2 **Identify the Fibre Channel (FC) port on the initiator.**

```
initiator# fcinfo hba-port
HBA Port WWN: 210000e08b195dae
 Port Mode: Initiator
 Port ID: 0
 OS Device Name: /dev/cfg/c8
 Manufacturer: QLogic Corp.
 Model: 375-3108-xx
 Firmware Version: 03.03.28
 FCode/BIOS Version: fcode: 1.13;
 Serial Number: not available
 Driver Name: qlc
 Driver Version: 20100408-3.01
 Type: unknown
 State: offline
 Supported Speeds: 1Gb 2Gb
 Current Speed: not established
 Node WWN: 200000e08b195dae
 NPIV Not Supported
.
.
.
```

- 3 **Create a host group.**
- 4 **Add the WWNs identified in the preceding output as members of the host group.**

```
target# stmfadm create-hg host-a 210000e08b195dae 210100e08b395dae
```

```
target# stmfadm add-hg-member -g host-a
```

- 5 **Create a target group.**
- 6 **Specify the target group members by adding the target names.**

```
target# stmfadm create-tg targets-0
```

Each SCSI target can be a member of only one target group.

```
target# stmfadm add-tg-member -g targets-0 wwn.guid-number
```

**7 Identify the GUID number for the LU.**

```
target# stmfadm list-lu -v
```

**8 Make the LU available by adding a view entry, specifying the host group name and the LU GUID number.**

```
target# stmfadm add-view -h host-a -t targets-0 -n 1 guid-number
```

## Configuring Fibre Channel Devices With COMSTAR

You can set up and configure a COMSTAR target on a SPARC system or x86 system in a FC network environment and make it accessible to the storage network. Ensure the following prerequisite step has been completed:

- [“How to Create a Logical Unit” on page 240](#)

## Configuring Fibre Channel Ports With COMSTAR

The Fibre Channel (FC) port provider can use the same HBAs that are used for the FC initiators. A given FC port can be used as an initiator or as a target, but not as both. You can also configure a dual port or quad port FC HBA so that a subset of the HBA ports is in target mode and the rest of the ports are in initiator mode.

The procedures in this section are specific to QLogic HBAs. The driver to use with HBAs in initiator mode is `qlc`, which is a QLogic driver. The `qlc` driver works only in initiator mode, and cannot be used for target mode. The COMSTAR driver for a target QLogic 4G HBA is `qlt`.

Because initiator mode and target mode use different drivers, the driver you attach to an HBA port defines its function as a target or initiator. You can specify a driver for all the ports by specifying the PCI device ID of the HBA. Or, you can configure the driver on a port-by-port basis. Both methods use the `update_drv` command and are described in this section. For more information, see [update\\_drv\(1M\)](#).

### ▼ How to Display Existing FC Port Bindings

Before making changes to the HBA ports, first check the existing port bindings.

**1 Become an administrator.****2 Display what is currently bound to the port drivers.**

In this example, the current binding is `pciex1077,2432`.

```
mdb -k
Loading modules: [unix krtld genunix specs ...
> ::devbindings -q qlc
```

```
30001617a08 pciex1077,2432, instance #0 (driver name: qlc)
300016177e0 pciex1077,2432, instance #1 (driver name: qlc)
> $q
```

## ▼ How to Set All FC Ports to a Specific Mode

This procedure changes all ports on all HBAs with a specific PCI device ID to target mode. The PCI device ID number binds a driver to a port, thereby setting all the HBA ports with that PCI device ID, for example, all QLogic 4G PCI express HBAs, to target mode.

### 1 Become an administrator.

### 2 Remove the current binding.

In this example, the `qlc` driver is actively bound to `pciex1077,2432`. You must remove the existing binding for `qlc` before you can add that binding to a new driver. Single quotation marks are required in this syntax.

```
update_drv -d -i 'pciex1077,2432' qlc
Cannot unload module: qlc
Will be unloaded upon reboot.
```

This message does *not* indicate an error. The configuration files have been updated but the `qlc` driver remains bound to the port until the system is rebooted.

### 3 Establish the new binding.

In this example, `qlt` is updated. Single quotes are required in this syntax.

```
update_drv -a -i 'pciex1077,2432' qlt
devfsadm: driver failed to attach: qlt
Warning: Driver (qlt) successfully added to system but failed to attach
```

This message does not indicate an error. The `qlc` driver remains bound to the port, until the system is rebooted. The `qlt` driver attaches when the system is rebooted.

### 4 Reboot the system to attach the new driver. Then, recheck the bindings.

```
init 6
.
.
.
mdb -k

Loading modules: [unix krtld genunix specfs dtrace ...
> ::devbindings -q qlt
30001615a08 pciex1077,2432, instance #0 (driver name: qlt)
30001615e0 pciex1077,2432, instance #1 (driver name: qlt)
> $q
```

### 5 Verify that the target mode framework has access to the HBA ports.

```
stmfadm list-target -v
Target: wwn.210100E08BA54E60
Operational Status : Offline
Provider Name : qlt(1)
```

```

Alias : -
Sessions : 0
Target: wwn.210100E08BA54E60
Operational Status : Offline
Provider Name : qlt(0)
Alias : -
Sessions : 0

```

## ▼ How to Set Selected FC Ports to Initiator or Target Mode

This procedure uses *path-based bindings*. It shows you how to use a specific device path to bind a port to a driver that is different from the driver to which it is currently bound.

### 1 Become an administrator.

### 2 Display a list of the HBA ports and their respective device paths.

This example shows the device paths for a single HBA with two ports.

```

luxadm -e port
/devices/pci@780/QLGC,qlc@0,1/fp@0,0:devctl CONNECTED
/devices/pci@780/QLGC,qlc@0/fp@0,0:devctl CONNECTED

```

### 3 Set the top port to target mode, and leave the bottom port in initiator mode.

Remove the initial `/devices` portion of the path, and include everything up to `/fp@0...`. The path with the `/devices` portion removed is the path to which the system binds the `qlt` driver.

Single quotation marks are required in this syntax.

```

update_drv -a -i '/pci@780/QLGC,qlc@0,1' qlt
devfsadm: driver failed to attach: qlt
Warning: Driver (qlt) successfully added to system but failed to attach.

```

This message does not indicate an error. The `qlc` driver remains bound to the port until reboot. The `qlt` driver attaches during reboot.

### 4 Reboot the system to attach the new driver. Then, recheck the bindings.

You should see that the port changed from initiator mode (`qlc`) to target mode (`qlt`).

```

init 6
.
.
.
mdb -k
.
.
.
> $q

```

### 5 Verify that the target mode framework has access to the HBA ports.

```

stmfadm list-target -v

```

## Making Logical Units Available for FC and FCoE

Simply registering a logical unit (LU) with the STMF framework does not make it available to hosts (initiators) on the network. You must make logical units visible to initiator hosts for Fibre Channel and FCoE configurations by mapping the logical unit. To determine which method to use and how to map the logical unit, see [“How to Make a Logical Unit Available to All Systems” on page 246](#). Both methods use the `stmfadm` command. The additional steps below are for FC and FCoE configurations.

### ▼ How to Make Logical Units Available for FC and FCoE

This procedure makes the LU available to all hosts or selected hosts for FC or FCoE configurations on a storage network. The steps are run on the host.

#### 1 Become an administrator.

#### 2 Make a LU available to hosts.

Obtain the global unique identification (GUID) number for the LU.

```
sbdadm list-lu
stmfadm list-lu -v
```

Identify the WWNs for the FC or FCoE ports of the host.

```
fcinfo hba-port
HBA Port WWN: *210000e08b83378d*
OS Device Name: /dev/cfg/c4
Manufacturer: Qlogic Corp.
Model: QLA2462
Firmware Version: 4.0.27
Fcode/BIOS Version: N/A
Type: N-port
State: online
Supported Speeds: 1Gb 2Gb 4Gb
Current Speed: 4Gb
Node WWN: 210000e08b83378d
HBA Port WWN: *210100e08ba3378d*
OS Device Name: /dev/cfg/c5
Manufacturer: Qlogic Corp.
Model: QLA2462
Firmware Version: 4.0.27
Fcode/BIOS Version: N/A
Type: N-port
State: online
Supported Speeds: 1Gb 2Gb 4Gb
Current Speed: 4Gb
Node WWN: 210100e08ba3378d
```

#### 3 Add a view and perform mapping.

Follow the instructions in [“How to Make a Logical Unit Available to All Systems” on page 246](#).

**4 Verify that the LU is visible on an Oracle Solaris initiator host by running the following script.**

```
#!/bin/ksh
fcinfo hba-port |grep "^HBA" | awk '{print $4}' | while read ln
do
 fcinfo remote-port -p $ln -s >/dev/null 2>&1
done
```

All subsequent LUNs appear in the format output, because the script forces the initiator to touch all the LUNs through all the ports. If you do not see the LUNs, run the format command again. If you still do not see the LUNs, ensure that the service is enabled on the target by using the `svcs stmf` command. Also ensure that you added view entries for the LU, as described in [“How to Make a Logical Unit Available to All Systems” on page 246](#).

**5 Verify that the LU is visible on other systems.**

- For a Linux initiator host, verify that the LU is visible by running the utility provided by the HBA vendor. The utility scans for configuration changes.
- For a Windows initiator host, verify that the logical unit is visible by selecting Control Panel → Administrative Tools → Computer Management → Disk Management. Then, from the Action menu, choose Rescan Disks.

## Configuring FCoE Devices With COMSTAR

You can set up and configure a COMSTAR target in a Fibre Channel over Ethernet (FCoE) network environment, then make it accessible to the storage network. Ensure that your system meets the necessary prerequisites before you begin:

- [“How to Create a Logical Unit” on page 240](#)
- [“How to Create an iSCSI Target” on page 241](#)
- Chapter 4, “Configuring Solaris iSCSI Initiators,” in *Oracle Solaris 11.1 Administration: SAN Configuration and Multipathing*

## Configuring FCoE Ports

- [“Enabling 802.3x PAUSE and Jumbo Frames on the Ethernet Interface” on page 253](#)
- [“How to Create FCoE Target Ports” on page 253](#)
- [“How to Verify That an FCoE Target Port Is Working” on page 254](#)
- [“How to Delete FCoE Target Ports” on page 254](#)

FCoE functionality is provided through Ethernet interfaces. Fibre Channel over Ethernet (FCoE) ports are logical entities associated with Ethernet interfaces. Within an Oracle Solaris system, a one-to-one mapping of FCoE ports and Ethernet interfaces exist. You can associate only one FCoE port with a given Ethernet interface. FCoE and IP cannot share the same Ethernet interface. So, before you create an FCoE port on an Ethernet interface, ensure that the interface has been unplumbed.

The FCoE port configuration is persistent across reboots. All configured FCoE ports are created and placed online automatically after the system reboots.

For FCoE target ports, you must enable the following service before creating FCoE target ports to get persistent behavior.

```
svcadm enable svc:/system/fcoe_target:default
```

## Enabling 802.3x PAUSE and Jumbo Frames on the Ethernet Interface

These settings can vary for different Ethernet hardware and drivers. In most cases, you must modify the `driver.conf` file of the Ethernet interface and then reboot. See the `driver.conf` file for your Ethernet interface for details on how to enable these features.

Perform one of the following prerequisite steps before you create an FCoE port on an Ethernet interface:

- Enable 802.3x (also called PAUSE) on the Ethernet interface.  
Doing so ensures a consistent Ethernet transport.
- Enable jumbo frames (> 2.5 Kbyte) on the Ethernet interface.  
A Fibre Channel data frame can be as large as 2136 bytes.

## ▼ How to Create FCoE Target Ports

### 1 Create an FCoE target port on a specified network interface.

```
fcadm create-fcoe-port -t nxge0
```

If the Ethernet interface you selected does not support multiple unicast address (for example, the VMware network interface), you are prompted to explicitly enable promiscuous mode on that interface.

### 2 Enable promiscuous mode, if prompted.

```
fcadm create-fcoe-port -t -f e1000g0
```

If no error messages appear, the FCoE target port is created, and the FCoE target is online.

## ▼ How to Verify That an FCoE Target Port Is Working

### 1 Display the FCoE ports that you created.

```
fcadm list-fcoe-ports
HBA Port WWN: 200000144fda7f66
Port Type: Target
MAC Name: nxge0
MTU Size: 9194
MAC Factory Address: 00144fda7f66
MAC Current Address: 0efc009a002a
Promiscuous Mode: On
```

### 2 Display all target mode Fibre Channel HBA ports on the host.

```
fcinfo hba-port -t
HBA Port WWN: 200000144fda7f66
Port Mode: Target
Port ID: 9a002a
OS Device Name: Not Applicable
Manufacturer: Sun Microsystems, Inc.
Model: FCoE Virtual FC HBA
Firmware Version: N/A
FCode/BIOS Version: N/A
Serial Number: N/A
Driver Name: COMSTAR FCOET
Driver Version: 1.0
Type: F-port
State: online
Supported Speeds: 1Gb 10 Gb
Current Speed: 10Gb
Node WWN: 100000144fda7f66
```

### 3 View a list of the FCoE target ports.

The `-v` option displays additional information about the target, along with SCSI session information for logged-in initiators.

```
stmfadm list-target -v
Target wwn.200000144FDA7F66
Operational Status: Online
Provider Name : fcoet
Alias : fcoet1
Sessions : 1
Initiator: wwn.210000E08B818343
Alias: #QLA2342 FW:v3.03.25 DVR:v8.02.14.01
Logged in since: ...
```

## ▼ How to Delete FCoE Target Ports

You can disable FCoE functionality when needed.

### 1 Take the FCoE target port offline.

```
stmfadm offline-target wwn.200000144fda7f66
```

## 2 Remove the FCoE target port.

```
fcadm delete-fcoe-port nxge0
```

# Configuring SRP Devices With COMSTAR

The SCSI RDMA Protocol accelerates the SCSI protocol by mapping the SCSI data transfer phases to Infiniband (IB) Remote Direct Memory Access (RDMA) operations. As a result, an SRP initiator can read and write data from a COMSTAR SRP target at high data rates with relatively low CPU utilization.

You can set up and configure a COMSTAR SRP target and make it available over an Infiniband (IB) fabric. The SRP target is available wherever a supported IB Host Channel Adapter (HCA) is installed on the target system.

- One SCSI target per IB HCA** – The COMSTAR SRP target uses a simple model in which each supported IB HCA is registered as a SCSI target. The SCSI target is a virtual object that contains a task router and acts as the connection between the SCSI transport (in this case, SRP) and the SCSI back end (STMF and SBD).

An HCA can contain multiple physical ports. The same SCSI target is shared between all the ports of an HCA. The SCSI target representing the HCA is automatically available for incoming connections through all the active ports of that HCA.

- SRP target eui identifier** – In the IB architecture, each HCA and each port is assigned a 64-bit GUID by the manufacturer. The COMSTAR SCSI target created for each HCA is given a name corresponding to the GUID of that HCA, which has the format: `eui.HCA-GUID`. For example, if the target system includes a supported IB HCA with an HCA GUID of `0003BA0001002E48`, then a SCSI target will be created using the name `eui.0003BA0001002E48`. The string `eui` stands for *extended unique identifier* and names a class of GUIDs that is used in both the SCSI and IB standards.
- SRP initiator eui identifier** – Similarly, SRP uses a 64-bit initiator GUID to identify the initiator system. The choice of which GUID to use is determined by the SRP initiator implementation. Many initiators use the GUID of the HCA that is being used for the outgoing connection. For example, an initiator using the GUID of `0003BA0001002EA5` is known to COMSTAR as `eui.0003BA0001002EA5`.

## Using COMSTAR Views With SRP

The COMSTAR view facility can be used to create target groups and host groups that restrict and configure which logical units (LUs) can be accessed through each SCSI target or initiator as described in [“Making SCSI Logical Units Available” on page 245](#). The `eui` identifier of the SRP initiator is added to a host group. The `eui` identifier of the SRP SCSI Target is added to a target group. The view entries for each LU then determine the particular set of LUs that each initiator can access.

## ▼ How to Enable the SRP Target Service

The COMSTAR port provider for the COMSTAR SRP target is managed by the Service Management Facility (SMF). The primary SRP target service is `svc:/system/ibsrp/target:default`, which can be abbreviated to `ibsrp/target`.

The SRP package is `storage/scsi-rdma/scsi-rdma-target`.

- 1 **Become an administrator.**
- 2 **Recursively enable the SRP target service.**  
`# svcadm enable -r ibsrp/target`
- 3 **Display the SRP target service information.**  
`# svcs -l ibsrp/target`

## ▼ How to Verify SRP Target Status

- 1 **Become an administrator.**
- 2 **Verify the presence of the expected SRP SCSI target on the system.**

```
srptadm list-target
Target HCA 21280001A0D0F0:
 Enabled : true
 SRP Target Name : eui.0021280001A0D0F0
 Operational Status : online
```

# Configuring and Managing the Oracle Solaris Internet Storage Name Service (iSNS)

---

This chapter provides an overview of the Internet Storage Name Service (iSNS), and describes how to configure the Oracle Solaris iSNS server, manage the iSNS server, and manage iSNS clients.

This is a list of the information in this chapter:

- [“The iSNS Technology \(Overview\)” on page 257](#)
- [“Configuring the iSNS Server” on page 259](#)
- [“Managing the iSNS Server and Clients” on page 265](#)

## The iSNS Technology (Overview)

The Internet Storage Name Service (iSNS) is a protocol that allows dynamic discovery of iSCSI initiators and targets within an IP storage area network SAN. The iSNS protocol enables identification, connection to, and management of iSCSI devices by providing the following services:

- **Name registration and discovery:** The source of data that is to be stored (known as the *initiator*) and the storage object (known as the *target*) register their attributes and address, and then can obtain information about accessible storage devices dynamically.
- **Discovery domains and logon control:** Resources in a typical storage network are divided into groups called *discovery domains*, which can be administered through network management applications. Discovery domains enhance security by providing access control to targets that are not enabled with their own access controls, while limiting the logon process of each initiator to a relevant subset of the available targets in the network.
- **State-change notification:** The iSNS server notifies relevant iSNS clients of network events, for example, a newly created disk Logical Unit Number (LUN), storage resources going offline, discovery domain membership changes and link failures in a network. These notifications let a network quickly adapt to changes in topology, which is key to scalability and availability. This is an optional service.

- **Entity status inquiry:** The iSNS server verifies that a iSNS client is available. As a result, a status change notification might be issued. This is an optional service.

In a simple configuration, the source of data that is to be stored (the initiator) exchanges data with a storage object (the target). The initiator can locate the target and the target always recognizes the initiator. For example, the Oracle StorageTek 5320 Network Attached Storage (NAS) appliance is a iSCSI target because it stores data. The data comes from various iSCSI clients such as a data management applications or network interface cards which act as initiators. However, in large and complex configurations, it is difficult and time-consuming to configure every initiator for every target and for every target to recognize every initiator. The iSNS server resolves this by using discovery and security mechanisms to dynamically and automatically identify initiators and targets, and manage their connections to authorized resources.

After a Oracle Solaris system has been configured as an iSNS server, all targets and initiators can register with the server. The targets and initiators become iSCSI *clients* or *nodes* of the iSNS server. These clients are members of the *default discovery domain*, the only domain in the *default discovery domain set*. When you enable the default discovery domain set, the iSNS server can provide the iSCSI Name Service (iSNS) for the clients in a simple manner.

To take advantage of the iSCSI Name Service's abilities, create several discovery domain sets and discovery domains. Then assign the clients to different domains, overlapping their memberships. The iSNS server keeps track of the clients' status as a member of one or more discovery domains. For example, when a new storage device is added to the storage network and is registered with the iSNS server, it is in the default discovery domain in the default discovery domain set. You then assign this target to the discovery domains whose initiators will use it as a resource. The iSNS server then removes this target as a member of the default discovery domain in the default discovery domain set.

All initiators and targets are assigned to at least one discovery domain. Assigning an initiator to one discovery domain restricts its access to those targets in the same discovery domain set. Assigning an initiator to several discovery domains allows it to find and use targets in all of the discovery domain sets that include the initiator's discovery domain. You can manage access to clients by disabling and enabling their discovery domain sets without affecting the clients in other discovery domain sets.

For example, a site has two discovery domain sets in addition to the default one: Production and Research. Within the two discovery domain sets are three domains in addition to the default one: Development, Operations, and Finance. The Development discovery domain is in the Research discovery domain set, Operations is in the Production domain set, and Finance is a member of both discovery domain sets. Each client has been assigned to the discovery domain set that uses it the most. A data application in the Operations discovery domain can locate and get access to storage devices in the Production discovery domain set because it is a member of that discovery domain set but it cannot get access to a storage device in the Research discovery domain set. A data application in the Finance discovery domain can locate storage devices in both the Production and Research discovery domain sets because it is a member of both sets. If

the Research discovery domain set were disabled, initiators in the Finance discovery domain would not have access to the Research storage devices but would continue to have access to those in the Production discovery domain set.

## Configuring the iSNS Server

You can configure the iSNS server using as described in the following task maps and sections.

Task	For Instructions
1. Accept the default properties of the iSNS server or change them.	
A. Notification of state changes of the server	<a href="#">“How to Set Notifications for Changes in Server State” on page 261</a>
B. Number of attempts to determine a client's availability	<a href="#">“How to Set the Number of Retries for Client Inquiries” on page 261</a>
C. Location of file that stores client data.	<a href="#">“How to Specify the Data Store Location” on page 261</a>
2. Enable the iSNS server and display the settings.	<a href="#">“How to Install the iSNS Server Package” on page 260</a> <a href="#">“How to Display the Current Server Configuration” on page 262</a>
3. Register all clients with the iSNS server.	Use the client's management interface's iSCSI configuration function to specify the IP address of the iSNS server and to allow discovery.
4. Enable the default discovery domain set.	<a href="#">“How to Enable the Default Discovery Domain Set” on page 263</a>

After these tasks, the iSNS server is operating in a minimal manner. All of the clients are in the default discovery domain and are unassigned. Each client can identify and get access to all of the other clients.

5. Create the discovery domain sets for your site.	<a href="#">“How to Create the Discovery Domain Sets” on page 263</a>
6. Create the discovery domains for your site.	<a href="#">“How to Create the Discovery Domains” on page 263</a>
7. Add each discovery domain to one or more discovery domain set.	<a href="#">“How to Add a Discovery Domain to a Discovery Domain Set” on page 264</a>
8. Assign clients to one or more discovery domains.	<a href="#">“How to Assign Clients to a Discovery Domain” on page 264</a>

9. Verify the membership of clients in discovery domains and the membership of discovery domains in discovery domain sets.	<a href="#">“How to Display the Status of a Discovery Domain Set” on page 266</a> <a href="#">“How to Display the Status of a Discovery Domain” on page 266</a> <a href="#">“How to Display the Status of Clients” on page 266</a>
----------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

The next section provides instructions for setting up the iSNS environment. The following topics are discussed:

- [“Setting Up the iSNS Administrative Settings” on page 260](#)
- [“Using the Command Line Interface to Configure iSNS” on page 262](#)

## Setting Up the iSNS Administrative Settings

This section provides the procedures for changing the default administrative settings of the iSNS service and for starting the iSNS daemon. If you change a setting after the iSNS server has been started, you need to refresh the iSNS server. If you change the data store location, you need to restart the iSNS server.

The following tasks are described in this section:

- [“How to Install the iSNS Server Package” on page 260](#)
- [“How to Set Notifications for Changes in Server State” on page 261](#)
- [“How to Set the Number of Retries for Client Inquiries” on page 261](#)
- [“How to Specify the Data Store Location” on page 261](#)

See the man page for the `isns(1M)` command details about these operations.

### ▼ How to Install the iSNS Server Package

Install the iSNS server package and start the iSNS service.

#### 1 Become an administrator.

#### 2 Install the iSNS server package.

```
pkg install service/storage/isns
```

#### 3 Enable the iSNS service.

```
svcadm enable isns_server
```

#### 4 Verify that the service is running.

```
svcs svc:/network/isns_server:default
STATE STIME FMRI
online 16:10:49 svc:/network/isns_server:default
```

## ▼ How to Set Notifications for Changes in Server State

By default, all clients are notified when the iSNS server is not available. To disable these notifications, change the *Management\_SCNs\_Enabled* property.

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Use the `svccfg` command to disable the property:

```
svccfg -s svc:/network/isns_server setprop config/Management_SCN_Enabled=false
```

- 3 Reload the server configuration:

```
svcadm refresh svc:/network/isns_server
```

## ▼ How to Set the Number of Retries for Client Inquiries

The default number of retries is 3. If the server does not get a response to three inquiries, it registers that client as unavailable. To change the number of retries, change the value of the *ESI Retry Threshold* property.

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Use the `svccfg` command to change the property to, for example, 6 retries:

```
svccfg -s svc:/network/isns_server setprop config/ESI_retry_threshold_count=6
```

- 3 Reload the server configuration:

```
svcadm refresh svc:/network/isns_server
```

## ▼ How to Specify the Data Store Location

The default location and name for the file that contains the client data is `/etc/isns/isnsdata.xml`. If you have a complex network environment that includes one or more backup iSNS servers, the data store must reside in a common location so that all servers can use it. Use the *data\_store\_location* property to specify the new location. You can also change the name of the file.

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Use the `svccfg` command to change the location to, for example, `/etc/isns2/isns_data.xml`:

```
svccfg -s svc:/network/isns_server setprop config/data_store_location="/etc/isns2/isns_data.xml"
```

- 3 If you change the data store location after the server has been enabled, you must restart the server:

```
svcadm restart svc:/network/isns_server
```

## Using the Command Line Interface to Configure iSNS

This section provides the procedures for configuring the iSNS server using the command line interface.

The following tasks are described in this section:

- “How to Display the Current Server Configuration” on page 262
- “How to Enable the Default Discovery Domain Set” on page 263
- “How to Create the Discovery Domain Sets” on page 263
- “How to Create the Discovery Domains” on page 263
- “How to Add a Discovery Domain to a Discovery Domain Set” on page 264
- “How to Assign Clients to a Discovery Domain” on page 264

These procedures use the `isnsadm(1M)` command. See the man page for a complete description of all of the command options.

### ▼ How to Display the Current Server Configuration

- The following command shows the properties of the iSNS server:

```
isnsadm show-config
Data Store Location: /etc/isns/isnsdata.xml
Entity Status Inquiry Non-Response Threshold: 3
Management SCN Enabled: yes
Authorized Control Node Names: -
```

## ▼ How to Enable the Default Discovery Domain Set

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Activate the default discovery domain set:

```
isnsadm enable-dd-set Default
```

## ▼ How to Create the Discovery Domain Sets

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Create a discovery domain set:

```
isnsadm create-dd-set set_name
```

- 3 Enable the discovery domain set:

```
isnsadm enable-dd-set set_name
```

- 4 View all the discovery domain sets, including the new one:

```
isnsadm list-dd-set -v
 DD Set name: Default
 State: Enabled
 DD Set name: set_name
 State: Enabled
```

The list of discovery domain sets includes the default discovery domain set as well as the new one.

## ▼ How to Create the Discovery Domains

New discovery domains are members of the default discovery domain set. After you create them, you add them to the new discovery domain set.

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

**2 Create the discovery domain:**

```
isnsadm create-dd domain_name
```

**3 View the new discovery domain in the Default discovery domain set:**

```
isnsadm list-dd-set
 DD name: name
 DD set(s): Default
```

**4 Create other discovery domains.****▼ How to Add a Discovery Domain to a Discovery Domain Set**

This task removes the discovery domain from the default discovery domain set and adds it the discovery domain set that you specify. Because the new discovery domain set has been enabled, all the clients in its discovery domains can be discovered by the iSNS server.

You do not need to have privileges to list the members of the discovery domains and discovery domain sets.

**1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.**

Roles contain authorizations and privileged commands. For more information about roles, see [“Initially Configuring RBAC \(Task Map\)”](#) in *Oracle Solaris 11.1 Administration: Security Services*.

**2 List the discovery domains to identify the one you want to add.**

```
isnsadm list-dd -v
```

**3 List the discovery domain sets to identify the one you want as the container for the new discovery domain.**

```
isnsadm list-dd-set
```

**4 Move the discovery domain to the discovery domain set that you want:**

```
isnsadm add-dd domain_name -s set_name
```

**5 View the new addition to the discovery domain set:**

```
isnsadm list-dd-set -v domain_name
```

**▼ How to Assign Clients to a Discovery Domain****Before You Begin**

Use the client's management interface to register the client. Using the iSCSI configuration function, specify the IP address of the iSNS server and allow discovery of the client by the iSNS server.

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Verify that the client has been registered with the iSNS server:

```
isnsadm list-node
 iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.45173FEA.engr
 Alias: STK5320_NAS
 Type: Target
 .
 iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.acct
 Alias:
 Type: Initiator
```

The output shows the clients' iSCSI names.

- 3 Verify the name of the discovery domain:

```
isnsadm list-dd
```

- 4 Add the client to the discovery domain:

```
isnsadm add-node -d domain_name iSCSI_Name
```

For example, to add the target called “STK5320\_NAS” to the Eng-dd discovery domain:

```
isnsadm add-node -d Eng-dd iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.engr
```

- 5 List all the clients in the discovery domain to verify the client has been added:

```
isnsadm list-dd -v domain_name
```

For example, to check the Eng-dd discovery domain:

```
isnsadm list-dd -v Eng-dd
 DD name: Eng-dd
 DD set: Development-dds
 iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.45173FEA.engr
 iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.acct
 iSCSI name: iqn.1986-03.com.sun:01:e00000000000.46fd8e2b
```

## Managing the iSNS Server and Clients

This section describes how to maintain the iSNS discovery domain sets and their members, the initiators and targets.

As the site grows, continue to add clients, discovery domains, and discovery domain sets as described in the following sections:

- “How to Create the Discovery Domain Sets” on page 263
- “How to Create the Discovery Domains” on page 263
- “How to Add a Discovery Domain to a Discovery Domain Set” on page 264
- “How to Assign Clients to a Discovery Domain” on page 264

This section provides the other procedures for managing the iSNS server, using the command line interface.

The following tasks are described in this section:

- “How to Display the Status of a Discovery Domain Set” on page 266
- “How to Display the Status of a Discovery Domain” on page 266
- “How to Display the Status of Clients” on page 266
- “How to Remove a Client from a Discovery Domain” on page 267
- “How to Remove a Discovery Domain from a Discovery Domain Set” on page 267
- “How to Disable a Discovery Domain Set” on page 268
- “How to Remove a Discovery Domain Set” on page 268

## ▼ How to Display the Status of a Discovery Domain Set

- Show the status of the discovery domain set and list the discovery domains that are its members:

```
isnsadm list-dd-set -v set_name
```

## ▼ How to Display the Status of a Discovery Domain

- Show the status of the discovery domain and lists the clients that are its members:

```
isnsadm list-dd -v domain_name
```

## ▼ How to Display the Status of Clients

- Select one of the following to display client status:

- Show the status of all clients:

```
isnsadm list-node -v
```

- Show the status of only the clients that are targets, that is, storage objects:

```
isnsadm list-node -t
```

## ▼ How to Remove a Client from a Discovery Domain

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 List the clients to identify the one you want to remove.

```
isnsadm list-node -v
iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.45173FEA.engr
Alias: STK5320_NAS
Type: Target
Network Entity: SE5310
Portal: 172.20.57.95:3260
 Portal Group: 1
Portal: 172.20.56.95:3260
 Portal Group: 1
DD Name: Research,Finance
```

The output shows the client's iSCSI name and the name of the discovery domains of which it is a member.

- 3 Remove the client from the discovery domain.

```
isnsadm remove-node -d domain_name iSCSI_name
```

## ▼ How to Remove a Discovery Domain from a Discovery Domain Set

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 List the discovery domains to identify the one you want to remove.

```
isnsadm list-dd -v
```

- 3 Remove the discovery domain from the discovery domain set.

```
isnsadm remove-dd set_name domain_name
```

## ▼ How to Disable a Discovery Domain Set

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 Deactivate a discovery domain set:

```
isnsadm disable-dd-set set_name
```

- 3 Verify that the state of the discovery domain set has changed to Disabled:

```
isnsadm list-dd-set set_name
```

## ▼ How to Remove a Discovery Domain Set

After you remove a discovery domain set, its discovery domains remain. A discovery domain must be a member of at least one discovery domain set.

- 1 Use the “iSNS Server Management” RBAC profile to obtain the authorizations needed for managing the iSNS service.

Roles contain authorizations and privileged commands. For more information about roles, see “Initially Configuring RBAC (Task Map)” in *Oracle Solaris 11.1 Administration: Security Services*.

- 2 List the discovery domain sets to identify the one you want to remove.

```
isnsadm list-dd-set -v
```

- 3 Remove the discovery domain set.

```
isnsadm remove-dd-set set_name
```

## The format Utility (Reference)

---

This chapter describes the format utility's menus and commands.

This is a list of the reference information in this chapter.

- “Recommendations and Requirements for Using the format Utility” on page 269
- “format Menu and Command Descriptions” on page 270
- “Rules for Input to format Commands” on page 276
- “Getting Help on the format Utility” on page 277

For a overview of when to use the format utility, see “format Utility” on page 168.

### Recommendations and Requirements for Using the format Utility

You must assume the root role or become an administrator to use the format utility. See “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*, Otherwise, the following error message is displayed when you try to use the format utility:

```
$ format
Searching for disks...done
No permission (or no disks found)!
```

# format Menu and Command Descriptions

The format menu contents are displayed as follows:

```

FORMAT MENU:
 disk - select a disk
 type - select (define) a disk type
 partition - select (define) a partition table
 current - describe the current disk
 format - format and analyze the disk
 fdisk - run the fdisk program (x86 only)
 repair - repair a defective sector
 label - write label to the disk
 analyze - surface analysis
 defect - defect list management
 backup - search for backup labels
 verify - read and display labels
 save - save new disk/partition definitions
 inquiry - show vendor, product and revision
 volname - set 8-character volume name
 !<cmd> - execute <cmd>, then return
 quit
format>

```

The following table describes the main menu items for the format utility.

**TABLE 13-1** The Main Menu Item Descriptions for the format Utility

Menu Item	Command or Menu?	Description
disk	Command	Lists all of the system's drives. Also lets you choose the disk you want to use in subsequent operations. This disk is referred to as the current disk.
type	Command	Identifies the manufacturer and model of the current disk. Also displays a list of known drive types. Choose the Auto configure option for all SCSI-2 disk drives.
partition	Menu	Creates and modifies slices. For more information, see <a href="#">“partition Menu” on page 272</a> .
current	Command	Displays the following information about the current disk: <ul style="list-style-type: none"> <li>■ Device name and device type</li> <li>■ Number of cylinders, alternate cylinders, heads and sectors</li> <li>■ Physical device name</li> </ul>

TABLE 13-1 The Main Menu Item Descriptions for the format Utility (Continued)

Menu Item	Command or Menu?	Description
format	Command	<p>Formats the current disk by using one of these sources of information in this order:</p> <ol style="list-style-type: none"> <li>1. Information that is found in the <code>format.dat</code> file</li> <li>2. Information from the automatic configuration process</li> <li>3. Information that you type at the prompt if no <code>format.dat</code> entry exists</li> </ol> <p>This command does not apply to IDE disks. IDE disks are preformatted by the manufacturer.</p>
fdisk	Menu	<p>x86 platform only: Runs the <code>fdisk</code> program to create a Solaris <code>fdisk</code> partition.</p> <p>The <code>fdisk</code> command cannot be used on disks with an EFI label that are greater than 1 terabyte in size.</p>
repair	Command	Repairs a specific block on the current disk.
label	Command	Writes a new label to the current disk.
analyze	Menu	Runs read, write, and compare tests. For more information, see “ <a href="#">analyze Menu</a> ” on page 273.
defect	Menu	Retrieves and displays defect lists. For more information, see “ <a href="#">defect Menu</a> ” on page 275. This feature does not apply to IDE disks. IDE disks manage defects automatically.
backup	Command	<p><b>VTOC</b> – Searches for backup labels.</p> <p><b>EFI</b> – Not supported.</p>
verify	Command	<p>Displays the following information about the current disk:</p> <ul style="list-style-type: none"> <li>■ Device name and device type</li> <li>■ Number of cylinders, alternate cylinders, heads and sectors</li> <li>■ Partition table</li> </ul>
save	Command	<p><b>VTOC</b> – Saves new disk and partition information.</p> <p><b>EFI</b> – Not applicable.</p>
inquiry	Command	<b>SCSI disks only</b> – Displays the vendor, product name, and revision level of the current drive.
volname	Command	Labels the disk with a new eight-character volume name that you specify.
quit	Command	Exits the format menu.

## partition Menu

The partition menu contents are displayed as follows:

```
format> partition
PARTITION MENU:
 0 - change '0' partition
 1 - change '1' partition
 2 - change '2' partition
 3 - change '3' partition
 4 - change '4' partition
 5 - change '5' partition
 6 - change '6' partition
 7 - change '7' partition
select - select a predefined table
modify - modify a predefined partition table
name - name the current table
print - display the current table
label - write partition map and label to the disk
quit
```

partition>

The following table describes the partition menu items.

**TABLE 13-2** Descriptions for partition Menu Items

Subcommand	Description
change 'n' partition	Enables you to specify the following information for the new partition: <ul style="list-style-type: none"> <li>■ Identification tag</li> <li>■ Permission flags</li> <li>■ Starting cylinder</li> <li>■ Size</li> </ul>
select	Enables you to choose a predefined partition table.
modify	Enables you to change all the slices in the partition table. This command is preferred over the individual change 'x' partition commands.
name	Enables you to specify a name for the current partition table.
print	Displays the current partition table.
label	Writes the partition map and the label to the current disk.
quit	Exits the partition menu.

## x86: fdisk Menu

The fdisk menu is available on x86 based systems only and appears similar to the following.

```

format> fdisk
 Total disk size is 8924 cylinders
 Cylinder size is 16065 (512 byte) blocks

 Cylinders
 Start End Length %
 ===== ===== ===== =====
 1 EFI 0 8924 8925 100

SELECT ONE OF THE FOLLOWING:
 1. Create a partition
 2. Specify the active partition
 3. Delete a partition
 4. Change between Solaris and Solaris2 Partition IDs
 5. Edit/View extended partitions
 6. Exit (update disk configuration and exit)
 7. Cancel (exit without updating disk configuration)
Enter Selection:

```

The following table describes the fdisk menu items.

**TABLE 13-3** x86: Descriptions for fdisk Menu Items

Menu Item	Description
Create a partition	Creates an fdisk partition. You must create a separate partition for each OS such as Oracle Solaris or DOS. There is a maximum of four partitions per disk. You are prompted for the size of the fdisk partition as a percentage of the disk.
Specify the active partition	Enables you to specify the partition to be used for booting. This menu item identifies where the first stage boot program looks for the second stage boot program.
Delete a partition	Deletes a previously created partition. This command destroys all the data in the partition.
Change between Solaris and Solaris2 Partition IDs	Changes partition IDs from 130 (0x82) to 191 (0xbf) and back again.
Edit/View extended partitions	Manages partition information that is generally used for booting.
Exit (update disk configuration and exit)	Writes a new version of the partition table and exits the fdisk menu.
Cancel (exit without updating disk configuration)	Exits the fdisk menu without modifying the partition table.

## analyze Menu

The analyze menu contents are displayed as follows:

```
format> analyze
```

```
ANALYZE MENU:
```

```
 read - read only test (doesn't harm SunOS)
 refresh - read then write (doesn't harm data)
 test - pattern testing (doesn't harm data)
 write - write then read (corrupts data)
 compare - write, read, compare (corrupts data)
 purge - write, read, write (corrupts data)
 verify - write entire disk, then verify (corrupts data)
 print - display data buffer
 setup - set analysis parameters
 config - show analysis parameters
 quit
analyze>
```

The following table describes the analyze menu items.

**TABLE 13-4** Descriptions for analyze Menu Items

Subcommand	Description
read	Reads each sector on the current disk. Repairs defective blocks as a default.
refresh	Reads then writes data on the current disk without harming the data. Repairs defective blocks as a default.
test	Writes a set of patterns to the disk without harming the data. Repairs defective blocks as a default.
write	Writes a set of patterns to the disk then reads back the data on the disk. Destroys existing data on the disk. Repairs defective blocks as a default.
compare	Writes a set of patterns to the disk, reads back the data, and then compares it to the data in the write buffer. Destroys existing data on the disk. Repairs defective blocks as a default.
purge	Removes all data from the disk so that the data cannot be retrieved by any means. Data is removed by writing three distinct patterns over the entire disk (or a section of the disk). If the verification passes, a hex-bit pattern is written over the entire disk (or a section of the disk).  Repairs defective blocks as a default.
verify	In the first pass, writes unique data to each block on the entire disk. In the next pass, reads and verifies the data. Destroys existing data on the disk. Repairs defective blocks as a default.
print	Displays the data in the read/write buffer.

TABLE 13-4 Descriptions for analyze Menu Items (Continued)

Subcommand	Description
setup	Enables you to specify the following analysis parameters: Analyze entire disk? yes Starting block number: <i>depends on drive</i> Ending block number: <i>depends on drive</i> Loop continuously? no Number of passes: 2 Repair defective blocks? yes Stop after first error? no Use random bit patterns? no Number of blocks per transfer: 126 (0/n/nn) Verify media after formatting? yes Enable extended messages? no Restore defect list? yes Restore disk label? yes
config	Displays the current analysis parameters.
quit	Exits the analyze menu.

## defect Menu

The defect menu contents are displayed as follows:

```
format> defect
DEFECT MENU:
 primary - extract manufacturer's defect list
 grown - extract manufacturer's and repaired defects lists
 both - extract both primary and grown defects lists
 print - display working list
 dump - dump working list to file
 quit
defect>
```

The following table describes the defect menu items.

TABLE 13-5 The defect Menu Item Descriptions

Subcommand	Description
primary	Reads the manufacturer's defect list from the disk drive and updates the in-memory defect list.
grown	Reads the grown defect list and then updates the in-memory defect list. Grown defects are defects that have been detected during analysis.
both	Reads both the manufacturer's defect list and the grown defect list. Then, updates the in-memory defect list.

TABLE 13-5 The defect Menu Item Descriptions (Continued)

Subcommand	Description
print	Displays the in-memory defect list.
dump	Saves the in-memory defect list to a file.
quit	Exits the defect menu.

## Rules for Input to format Commands

When you use the `format` utility, you need to provide various kinds of information. This section describes the rules for this information. For information on using `format`'s help facility when you specify data, see [“Getting Help on the format Utility” on page 277](#).

### Specifying Numbers to format Commands

Several places in the `format` utility require number as input. You must either specify the appropriate data or select a number from a list of choices. In either case, the help facility causes `format` to display the upper and lower limits of the number expected. Simply enter the appropriate number. The number is assumed to be in decimal format unless a base is explicitly specified as part of the number (for example, `0x` for hexadecimal).

The following are examples of integer input:

```
Enter number of passes [2]: 34
Enter number of passes [34] 0xf
```

### Specifying format Command Names

Command names are needed as input whenever the `format` utility displays a menu prompt. You can abbreviate the command names, as long as what you type is sufficient to uniquely identify the command desired.

For example, use `p` to access the `partition` menu from the `format` menu. Then, type `p` to display the current slice table.

```
format> p
PARTITION MENU:
 0 - change '0' partition
 1 - change '1' partition
 2 - change '2' partition
 3 - change '3' partition
 4 - change '4' partition
 5 - change '5' partition
```

```

6 - change '6' partition
7 - change '7' partition
select - select a predefined table
modify - modify a predefined partition table
name - name the current table
print - display the current table
label - write partition map and label to the disk
quit
partition> p

```

## Specifying Disk Names to format Commands

At certain points in the `format` utility, you must name something. In these cases, you are free to specify any string you want for the name. If the name has white space in it, the entire name must be enclosed in double quotation marks (""). Otherwise, only the first word of the name is used.

For example, if you want to identify a specific partition table for a disk, you can use the `name` subcommand that is available from the `partition` menu:

```

partition> name
Enter table name (remember quotes): "new disk3"

```

## Getting Help on the format Utility

The `format` utility provides a help facility that you can use whenever the `format` utility is expecting input. You can request help about what input is expected by typing a question mark (?). The `format` utility displays a brief description of what type of input is needed.

If you type a ? at a menu prompt, a list of available commands is displayed.

The man pages associated with the `format` utility include the following:

- `format(1M)` – Describes the basic `format` utility capabilities and provides descriptions of all command-line variables.
- `format.dat(4)` – Describes disk drive configuration information for the `format` utility.



## Managing File Systems (Overview)

---

Managing file systems is one of your most important system administration tasks.

This is a list of the overview information in this chapter:

- “What's New in Oracle Solaris File Systems?” on page 279
- “Where to Find File System Management Tasks” on page 280
- “Overview of File Systems” on page 280
- “Default Oracle Solaris File Systems” on page 285
- “Overview of Mounting and Unmounting File Systems” on page 287
- “Determining a File System's Type” on page 291

## What's New in Oracle Solaris File Systems?

### Shared File Systems Across Boot Environments

A mechanism that enables automatic data sharing across different boot environments is available in this release. These shared directories are stored in the `rpool/VARSHARE` file system that is mounted at `/var/share`. Placement of shared data in the `/var` directory reduces the amount of space needed for all boot environments.

For example:

```
ls /var/share
audit cores crash mail
```

The `/var/share` file system generally requires no administration, with the exception of ensuring that the `/var` components do not fill the root file system.

Symbolic links are automatically created from `/var` to the `/var/share` components listed above for compatibility purposes. For more information, see [datasets\(5\)](#).

## Where to Find File System Management Tasks

Use these references to find step-by-step instructions for managing file systems.

File System Management Task	For More Information
Connect and configure new disk devices.	Chapter 7, “Managing Disks (Overview)”
Create and mount new file systems.	Chapter 15, “Creating and Mounting File Systems (Tasks)”
Make remote files available to users.	Chapter 2, “Network File System Administration (Tasks),” in <i>Managing Network File Systems in Oracle Solaris 11.1</i>

## Overview of File Systems

A file system is a structure of directories that is used to organize and store files.

The term *file system* is used to describe the following:

- A particular type of file system: disk-based, network-based, or virtual
- The entire file tree, beginning with the root (*/*) directory
- The data structure of a disk slice or other media storage device
- A portion of a file tree structure that is attached to a mount point on the main file tree so that the files are accessible

Usually, you know from the context which meaning is intended.

The Oracle Solaris OS uses the *virtual file system* (VFS) architecture, which provides a standard interface for different file system types. The VFS architecture enables the kernel to handle basic operations, such as reading, writing, and listing files. The VFS architecture also makes it easier to add new file systems.

## Types of Oracle Solaris File Systems

The Oracle Solaris OS supports three types of file systems:

- Disk-based
- Network-based
- Virtual

To identify the file system type, see “[Determining a File System's Type](#)” on page 291.

## Oracle Solaris Disk-Based File Systems

*Disk-based file systems* are stored on physical media such as hard disks and DVDs. Disk-based file systems can be written in different formats. The available formats are described in the following table.

Disk-Based File System	Format Description
ZFS	ZFS is the default disk-based and root file system in the Oracle Solaris 11 release. For more information, see the <a href="#">Oracle Solaris 11.1 Administration: ZFS File Systems</a> .
UFS	Legacy UNIX file system (based on the BSD Fat Fast File system that was provided in the 4.3 Tahoe release).
HSFS	High Sierra, Rock Ridge, and ISO 9660 file system. High Sierra is the first CD-ROM file system. ISO 9660 is the official standard version of the High Sierra file system. The HSFS file system is used on CD-ROMs, and is a read-only file system. Oracle Solaris HSFS supports Rock Ridge extensions to ISO 9660. When present on a CD-ROM, these extensions provide all file system features and file types, except for writability and hard links.
PCFS	PC file system, which allows read- and write- access to data and programs on DOS-formatted disks that are written for DOS-based personal computers.
UDFS	The Universal Disk Format (UDFS) file system, the industry-standard format for storing information on the optical media technology called DVD (Digital Versatile Disc or Digital Video Disc).

Each type of disk-based file system is customarily associated with a particular media device, as follows:

- ZFS or UFS with hard disk
- HSFS with CD-ROM
- PCFS with USB diskette
- UDF with DVD

However, these associations are not restrictive. For example, DVDs can have ZFS or UFS file systems created on them.

### The Universal Disk Format (UDFS) File System

For information about creating a UDFS file system on removable media, see [“How to Create a File System on Removable Media”](#) on page 23.

The UDF file system is the industry-standard format for storing information on *DVD* (Digital Versatile Disc or Digital Video Disc) optical media.

The UDF file system is provided as dynamically loadable 32-bit and 64-bit modules, with system administration utilities for creating, mounting, and checking the file system on both SPARC and x86 platforms. The Oracle Solaris UDF file system works with supported ATAPI and SCSI DVD drives, CD-ROM devices, and disk drives. In addition, the Oracle Solaris UDF file system is fully compliant with the UDF 1.50 specification.

The UDF file system provides the following features:

- Ability to access the industry-standard CD-ROM and DVD-ROM media when they contain a UDF file system
- Flexibility in exchanging information across platforms and operating systems
- A mechanism for implementing new applications rich in broadcast-quality video, high-quality sound, and interactivity using the DVD video specification based on UDF format

The following features are not included in the UDF file system:

- Support for write-once media, (CD-RW), with either the sequential disk-at-once recording and incremental recording

The UDF file system requires the following:

- Supported SPARC or x86 platform
- Supported CD-ROM or DVD-ROM device

The Oracle Solaris UDF file system implementation provides the following:

- Support for industry-standard read/write UDF version 1.50
- Fully internationalized file system utilities

## Network-Based File Systems

*Network-based file systems* can be accessed from the network. Typically, network-based file systems reside on one system, typically a server, and are accessed by other systems across the network.

With the NFS service, you can provide distributed *resources* (files or directories) by sharing them from a server and mounting them on individual clients. For more information, see [“The NFS Environment” on page 288](#).

With the Oracle SMB service, you can provide distributed *resources* (files or directories) to Windows and Mac OS systems by sharing them from a server and mounting them on individual clients. For more information, see [“The Oracle Solaris SMB Service” on page 290](#).

## Virtual File Systems

*Virtual file systems* are memory-based file systems that provide access to special kernel information and facilities. Most virtual file systems do not use file system disk space. Also, some virtual file systems, such as the temporary file system (TMPFS), use the swap space on a disk.

## Temporary File System

The temporary file system (TMPFS) uses local memory for file system reads and writes. Using TMPFS can improve system performance by saving the cost of reading and writing temporary files to a local disk or across the network. For example, temporary files are created when you compile a program. The OS generates a much disk activity or network activity while manipulating these files. Using TMPFS to hold these temporary files can significantly speed up their creation, manipulation, and deletion.

Files in TMPFS file systems are not permanent. These files are deleted when the file system is unmounted and when the system is shut down or rebooted.

TMPFS is the default file system type for the `/tmp` directory in the Oracle Solaris OS. You can copy or move files into or out of the `/tmp` directory, just as you would in a ZFS or UFS file system.

The TMPFS file system uses swap space as a temporary backing store.

If a system with a TMPFS file system does not have adequate swap space, two problems can occur:

- The TMPFS file system can run out of space, just as regular file systems do.
- Because TMPFS allocates swap space to save file data (if necessary), some programs might not execute because of insufficient swap space.

For information about creating TMPFS file systems, see [Chapter 15, “Creating and Mounting File Systems \(Tasks\)”](#). For information about increasing swap space, see [Chapter 16, “Configuring Additional Swap Space \(Tasks\)”](#).

## The Loopback File System

The loopback file system (LOFS) lets you create a new virtual file system so that you can access files by using an alternative path name. For example, you can create a loopback mount of the root (`/`) directory on `/tmp/newroot`. This loopback mounts make the entire file system hierarchy appear as if it is duplicated under `/tmp/newroot`, including any file systems mounted from NFS servers. All files will be accessible either with a path name starting from root (`/`), or with a path name that starts from `/tmp/newroot`.

For information on how to create LOFS file systems, see [Chapter 15, “Creating and Mounting File Systems \(Tasks\)”](#).

## Process File System

The process file system (PROCFS) resides in memory and contains a list of active processes, by process number, in the `/proc` directory. Information in the `/proc` directory is used by commands such as `ps`. Debuggers and other development tools can also access the address space of the processes by using file system calls.



---

**Caution** – Do not delete files in the `/proc` directory. The deletion of processes from the `/proc` directory does not kill them. `/proc` files do not use disk space, so there is no reason to delete files from this directory.

---

The `/proc` directory does not require administration.

## Additional Virtual File Systems

These additional types of virtual file systems are listed for your information. They do not require administration.

---

Virtual File System	Description
CTFS	CTFS (the contract file system) is the interface for creating, controlling, and observing contracts. A contract enhances the relationship between a process and the system resources it depends on by providing richer error reporting and (optionally) a means of delaying the removal of a resource.  The service management facility (SMF) uses process contracts (a type of contract) to track the processes which compose a service, so that a failure in a part of a multi-process service can be identified as a failure of that service.
FIFOFS (first-in first-out)	Named pipe files that give processes common access to data
FDFS (file descriptors)	Provides explicit names for opening files by using file descriptors
MNTFS	Provides read-only access to the table of mounted file systems for the local system
NAMEFS	Used mostly by STREAMS for dynamic mounts of file descriptors on top of files
OBJFS	The OBJFS (object) file system describes the state of all modules currently loaded by the kernel. This file system is used by debuggers to access information about kernel symbols without having to access the kernel directly.
SHAREFS	Provides read-only access to the table of shared file systems for the local system
SPECFS (special)	Provides access to character special devices and block devices
SWAPFS	Used by the kernel for swapping

---

## Extended File Attributes

The ZFS, UFS, NFS, and TMPFS file systems have been enhanced to include extended file attributes. Extended file attributes enable application developers to associate specific attributes to a file. For example, a developer of an application used to manage a windowing system might choose to associate a display icon with a file. Extended file attributes are logically represented as files within a hidden directory that is associated with the target file.

You can use the `runat` command to add attributes and execute shell commands in the extended attribute namespace. This namespace is a hidden attribute directory that is associated with the specified file.

To use the `runat` command to add attributes to a file, you first have to create the attributes file.

```
$ runat filea cp /tmp/attrdata attr.1
```

Then, use the `runat` command to list the attributes of the file.

```
$ runat filea ls -l
```

For more information, see the [runat\(1\)](#) man page.

Many Oracle Solaris file system commands have been modified to support file system attributes by providing an attribute-aware option. Use this option to query, copy, or find file attributes. For more information, see the specific man page for each file system command.

## Swap Space

The Oracle Solaris OS uses some disk slices for temporary storage rather than for file systems. These slices are called *swap* slices, or *swap space*. Swap space is used for virtual memory storage areas when the system does not have enough physical memory to handle current processes.

Since many applications rely on swap space, you should know how to plan for, monitor, and add more swap space, when needed. For an overview about swap space and instructions for adding swap space, see [Chapter 16, “Configuring Additional Swap Space \(Tasks\)”](#).

# Default Oracle Solaris File Systems

The ZFS file system is hierarchical, starting with the root directory (`/`) and continuing downwards through a number of directories. The Oracle Solaris installation process enables you to install a default set of directories and uses a set of conventions to group similar types of files together.

For a brief overview of Oracle Solaris file systems and directories, see [filesystem\(5\)](#).

The following table provides a summary of the default Oracle Solaris file systems.

TABLE 14-1 The Default Oracle Solaris File Systems

File System or Directory	File System Type	Description
root (/)	ZFS	The top of the hierarchical file tree. The root (/) directory contains the directories and files that are critical for system operation, such as the kernel, the device drivers, and the programs used to boot the system. The root (/) directory also contains the mount point directories where local and remote file systems can be attached to the file tree.
/usr	ZFS	System files and directories that can be shared with other users. Files that run only on certain types of systems are in the /usr directory (for example, SPARC executables). Files that can be used on all types of systems, such as the man pages, might be placed in the /usr/share directory.
/export/home or /home	NFS or ZFS	The mount point for user home directories, which store user work files. By default, the /home directory is an automounted file system.
/var	ZFS	System files and directories that are likely to change or grow over the life of the local system. These include system logs, such as vi and ex backup files.
/opt	NFS or ZFS	Optional mount point for third-party software. On some systems, the /opt directory might be a UFS file system or ZFS file system.
/tmp	TMPFS	Temporary files, which are removed each time the system is booted or the /tmp file system is unmounted.
/proc	PROCFS	A list of active processes, by process number.
/etc/mnttab	MNTFS	A virtual file system that provides read-only access to the table of mounted file systems for the local system.
/system/volatile	TMPFS	A memory-based file system for storing temporary files that are not needed after the system is booted.
/system/contract	CTFS	A virtual file system that maintains contract information.
/system/object	OBJFS	A virtual file system that is used by debuggers to access information about kernel symbols without having to access the kernel directly.

# Overview of Mounting and Unmounting File Systems

Before you can access the files on a file system, you need to mount the file system. When you mount a file system, you attach that file system to a directory (*mount point*) and make it available to the system. The root (/) file system is always mounted. Any other file system can be connected or disconnected from the root (/) file system.

Most file systems are automatically mounted by SMF services at system boot time. Generally, you do not need to mount or unmount file systems manually. For more information about mounting different file system types, see [“Mounting and Unmounting Oracle Solaris File Systems” on page 294](#).

When you mount a file system, any files or directories in the underlying mount point directory are unavailable as long as the file system is mounted. These files are not permanently affected by the mounting process. They become available again when the file system is unmounted. However, mount directories are typically empty because you usually do not want to obscure existing files.

For step-by-step instructions on how to mount file systems, see [“Mounting and Unmounting Oracle Solaris File Systems” on page 294](#).

## The Mounted File System Table

Whenever you mount or unmount a file system, the `/etc/mnttab` (mount table) file is modified with the list of currently mounted file systems. You can display the contents of this file by using the `cat` or `more` commands. However, you cannot edit this file. Here is an example of an `/etc/mnttab` file:

```
$ more /etc/mnttab
rpool/ROOT/zfsBE / zfs dev=3390002 0
/devices /devices devfs dev=8580000 1337114941
/dev /dev dev dev=85c0000 1337114941
ctfs /system/contract ctfs dev=8680001 1337114941
proc /proc proc dev=8600000 1337114941
mnttab /etc/mnttab mntfs dev=86c0001 1337114941
swap /system/volatile tmpfs xattr,dev=8700001 1337114941
objfs /system/object objfs dev=8740001 1337114941
sharefs /etc/dfs/sharetab sharefs dev=8780001 1337114941
/usr/lib/libc/libc_hwcapi2.so.1 /lib/libc.so.1 lofs dev=3390002 13371149
fd /dev/fd fd rw,dev=8880001 1337114969
rpool/ROOT/zfsBE/var /var zfs rw,devices, \
setuid,nonbmand,exec,
rstchown,xattr,atime,dev=3390003 1337114969
swap /tmp tmpfs xattr,dev=8700002 1337114969
rpool/VARSHARE /var/share zfs rw,devices,setuid,nonbmand,exec,
rstchown,xattr,atime,dev=3390004 1337114969
```

## The Virtual File System Table

Most file systems are mounted automatically by an SMF service at system boot time.

You might need to edit the `/etc/vfstab` file to mount legacy or remote file systems or to make changes to the ZFS swap volume. For information about changing a ZFS swap volume, see [Chapter 16, “Configuring Additional Swap Space \(Tasks\)”](#).

To add an entry for mounting a legacy or remote file system, the information you need to specify is as follows:

- The device or the NFS server where the file system resides
- The file system mount point
- File system type
- Whether you want the file system to mount automatically when the system boots (by using the `mountall` command)
- Any mount options

The following `vfstab` example is from a system that has a ZFS root file system. In addition, this system is mounting a remote file system, `/users/data`, from the NFS server, `neo`.

```
cat /etc/vfstab
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
fd - /dev/fd fd - no -
/proc - /proc proc - no -
/dev/zvol/dsk/rpool/swap - - swap - no -
/devices - /devices devfs - no -
sharefs - /etc/dfs/sharetabsharefs - no -
ctfs - /system/contract ctfs - no -
objfs - /system/object objfs - no -
swap - /tmp tmpfs - yes -
neo:/users/data - /data nfs - yes -
```

ZFS file systems are mounted automatically by the SMF service at boot time. You can mount ZFS file systems from the `vfstab` by using the legacy mount feature. For more information, see [Oracle Solaris 11.1 Administration: ZFS File Systems](#).

For descriptions of each `/etc/vfstab` field and information on how to edit and use the file, see [vfstab\(4\)](#).

## The NFS Environment

NFS is a distributed file system service that can be used to share *resources* (files or directories) from one system, typically a server, with other systems on the network. For example, you might want to share third-party applications or source files with users on other systems.

NFS makes the actual physical location of the resource irrelevant to the user. Instead of placing copies of commonly used files on every system, NFS allows you to place one copy on one system's disk and let all other systems access it from the network. Under NFS, remote files are virtually indistinguishable from local files.

For more information, see [Chapter 4, “Managing Network File Systems \(Overview\)”](#), in *Oracle Solaris Administration: Network Services*.

A system becomes an NFS server if it has resources to share on the network. A server keeps a list of currently shared resources and their access restrictions (such as read/write or read-only access).

When you share a resource, you make it available for mounting by remote systems.

You can share a resource in these ways:

- Create a ZFS share by setting the ZFS `share.nfs` property. For example:  

```
zfs set share.nfs=on tank/home
```
- Create a legacy share by using the `share` command.  

```
share -F nfs /ufsfs
```

For a complete description of NFS, see [Chapter 4, “Managing Network File Systems \(Overview\)”](#), in *Oracle Solaris Administration: Network Services*.

## NFS Version 4

Oracle's implementation of the NFS version 4 distributed file access protocol is included in the Oracle Solaris release.

NFS version 4 integrates file access, file locking, and mount protocols into a single, unified protocol to ease traversal through a firewall and improve security. The Oracle Solaris implementation of NFS version 4 is fully integrated with Kerberos V5, also known as SEAM, thus providing authentication, integrity, and privacy. NFS version 4 also enables the negotiation of security flavors to be used between the client and the server. With NFS version 4, a server can offer different security flavors for different file systems.

For more information about NFS Version 4 features, see [“What's New With the NFS Service”](#) in *Oracle Solaris Administration: Network Services*.

## Automounting (autofs)

You can mount NFS file system resources by using a client-side service called *automounting* (or *autofs*). The *autofs* service enables a system to automatically mount and unmount NFS resources whenever you access them. The resource remains mounted as long as you remain in the directory and are using a file within that directory. If the resource is not accessed for a certain period of time, it is automatically unmounted.

The autofs service provides the following features:

- NFS resources don't need to be mounted when the system boots, which saves booting time.
- Users don't need to know the root password to mount and unmount NFS resources.
- Network traffic might be reduced because NFS resources are mounted only when they are in use.

The autofs service is initialized by the automount utility, which runs automatically when a system is booted. The automountd daemon runs continuously and is responsible for the mounting and unmounting of NFS file systems on an as-needed basis. By default, the /home file system is mounted by the automount daemon.

With autofs, you can specify multiple servers to provide the same file system. This way, if one of these servers is down, autofs can try to mount the file system from another machine.

For complete information on how to set up and administer autofs, see [Chapter 2, “Network File System Administration \(Tasks\),”](#) in *Managing Network File Systems in Oracle Solaris 11.1*.

## The Oracle Solaris SMB Service

The Oracle Solaris OS provides a Server Message Block (SMB) protocol server and client implementation that includes support for numerous SMB dialects including NT LM 0.12 and Common Internet File System (CIFS). The terms CIFS and SMB can be considered interchangeable.

The Solaris SMB server allows a native Oracle Solaris system to serve files as SMB *shares* to SMB enabled clients that mount the file system shares. A Windows, Mac OS, or Solaris client can interoperate with the Solaris SMB server as it would with a Windows server. A Solaris SMB server can operate in either workgroup mode or in domain mode. In workgroup mode, the Solaris SMB server is responsible for authenticating users locally when access is requested to shared resources. This authentication process is referred to as local login. In domain mode, the Solaris SMB server uses pass-through authentication, in which user authentication is delegated to a domain controller.

For more information, see [Managing SMB File Sharing and Windows Interoperability in Oracle Solaris 11.1](#).

## Determining a File System's Type

You can determine a file system's type by using one of the following:

- The FS type field in the virtual file system table (the `/etc/vfstab` file)
- The `/etc/default/fs` file for local file systems
- The `/etc/dfs/fstypes` file for NFS file systems

## How to Determine a File System's Type

These commands work whether or not the file system is mounted.

If you have the raw device name of a disk slice, you can use the `fstyp` command to determine a file system's type (if the disk slice contains a file system). For more information, see [fstyp\(1M\)](#).

**EXAMPLE 14-1** Determining a File System's Type

The following example uses the `fstyp` command to determine the file system type.

```
fstyp /dev/rdisk/c0t0d0s0
zfs
```

The following example uses the `/etc/vfstab` file to determine the file system type for the `/legacy` file system.

```
$ grep /legacy /etc/vfstab
/dev/dsk/c0t3d0s6 /dev/rdisk/c0t3d0s6 /legacy ufs 2 yes -
```



# Creating and Mounting File Systems (Tasks)

---

This chapter describes how to create and mount ZFS, temporary (TMPFS), and loopback (LOFS) file systems. Because TMPFS and LOFS are virtual file systems, you actually “access” them by mounting them. In addition, creating and mounting a legacy UFS file system is also covered.

This is a list of the information in this chapter:

- [“Creating Oracle Solaris File Systems” on page 293](#)
- [“Mounting and Unmounting Oracle Solaris File Systems” on page 294](#)
- [“Creating and Mounting Oracle Solaris File Systems” on page 298](#)

## Creating Oracle Solaris File Systems

This section provides an overview of creating Oracle Solaris file systems.

### Creating ZFS File Systems

A ZFS file system is not directly tied to a specific disk partition. A ZFS file system is contained within a ZFS storage pool that can contain many devices. Either whole disks or disk slices can be added to a ZFS storage pool. Within a pool, you will probably want to create additional file systems. File systems provide points of administration that allow you to manage different sets of data within the same pool.

A ZFS file system is created by using the `zfs create` command. A ZFS file system is automatically mounted when it is created. For more information, see [“How to Create an ZFS File System” on page 298](#).

For more information about creating ZFS storage pools and file systems, see [Oracle Solaris 11.1 Administration: ZFS File Systems](#).

## Creating a Temporary File System

A *temporary file system (TMPFS)* uses local memory for file system reads and writes. TMPFS file systems can improve system performance by saving the cost of reading and writing temporary files to a local disk or across the network. Files in TMPFS file systems do not survive across reboots or unmounts.

If you create multiple TMPFS file systems, be aware that they all use the same system resources. Files created under one TMPFS file system use up space available for any other TMPFS file system, unless you limit TMPFS sizes by using the `-o size` option of the `mount` command.

For more information about creating and mounting a TMPFS file system, see [“How to Create and Mount a TMPFS File System”](#) on page 300.

For general information, see the `tmpfs(7FS)`.

## Creating a LOFS File System

An *LOFS file system* is a virtual file system that provides an alternate path to an existing file system. When other file systems are mounted onto an LOFS file system, the original file system does not change.

For more information about creating and mounting a LOFS file system, see [“How to Create and Mount an LOFS File System”](#) on page 301.

For general information, see the `lofs(7FS)`.

---

**Note** – Be careful when creating LOFS file systems. Because LOFS file systems are virtual file systems, the potential for confusing both users and applications is enormous.

---

## Mounting and Unmounting Oracle Solaris File Systems

ZFS file systems are mounted and unmounted automatically. You can make a legacy UFS file system available by mounting it, which attaches the file system to the system directory tree at the specified mount point. The root (`/`) file system is always mounted.

The following table provides guidelines on mounting file systems based on how you use them.

Mount Type Needed	Suggested Mount Method
Local or remote file systems that need to be mounted infrequently.	The mount command that you type manually from the command line.

Mount Type Needed	Suggested Mount Method
Local legacy UFS file systems that need to be mounted frequently. Local ZFS file systems are automatically mounted by an SMF service.	The <code>/etc/vfstab</code> file, which mounts the file system automatically when the system is booted in multi user state.
Remote legacy UFS file systems, such as home directories, that need to be mounted frequently.	<ul style="list-style-type: none"> <li>■ The <code>/etc/vfstab</code> file, which automatically mounts the file system when the system is booted in multiuser state.</li> <li>■ <code>autofs</code>, which automatically mounts the file system when you access it or unmounts the file system when you change to another directory.</li> </ul>

For more information on mounting removable media, see “[Managing Removable Media \(Overview\)](#)” on page 17.

You can determine which file systems are already mounted by using the `mount` command:

```
$ mount [-v]
```

The `-v` displays the list of mounted file systems in verbose mode.

#### EXAMPLE 15-1 Determining Which File Systems Are Mounted

This example shows how to use the `mount` command to display information about the file systems that are currently mounted.

```
$ mount
/ on rpool/ROOT/zfsBE read/write/setuid/devices/rstchown/dev=3390002 on Tue ...
/devices on /devices read/write/setuid/devices/rstchown/dev=8580000 on Tue May 15 ...
/dev on /dev read/write/setuid/devices/rstchown/dev=85c0000 on Tue May 15 14:49:01 2012
/system/contract on ctfs read/write/setuid/devices/rstchown/dev=8680001 on Tue May 15 ...
/proc on proc read/write/setuid/devices/rstchown/dev=8600000 on Tue May 15 14:49:01 2012
/etc/mnttab on mnttab read/write/setuid/devices/rstchown/dev=86c0001 on Tue May 15 14:49:01 ...
/system/volatile on swap read/write/setuid/devices/rstchown/xattr/dev=8700001 on Tue May 15 ...
/system/object on objfs read/write/setuid/devices/rstchown/dev=8740001 on Tue May 15 ...
/etc/dfs/sharetab on sharefs read/write/setuid/devices/rstchown/dev=8780001 on Tue May 15 ...
/lib/libc.so.1 on /usr/lib/libc/libc_hwcaps2.so.1 read/write/setuid/devices/rstchown/dev ...
/dev/fd on fd read/write/setuid/devices/rstchown/dev=8880001 on Tue May 15 14:49:29 2012
/var on rpool/ROOT/zfsBE/var read/write/setuid/devices/rstchown/nonbmand/exec/xattr/atime/ ...
/tmp on swap read/write/setuid/devices/rstchown/xattr/dev=8700002 on Tue May 15 14:49:29 2012
/var/share on rpool/VARSHARE read/write/setuid/devices/rstchown/nonbmand/exec/xattr/atime/ ...
/home/rimmer on pluto:/export/home/rimmer remote/read/write/setuid/xattr/...
```

This example shows how to use the `zfs mount` command to display information about ZFS file systems that are currently mounted.

```
$ zfs mount
rpool/ROOT/zfsBE /
rpool/ROOT/zfsBE /var
rpool/VARSHARE /var/share
```

**EXAMPLE 15-1** Determining Which File Systems Are Mounted (Continued)

```
tank/home /tank/home
```

## Field Descriptions for the /etc/vfstab File

An entry in the /etc/vfstab file has seven fields, which are described in the following table.

**TABLE 15-1** Field Descriptions for the /etc/vfstab File

Field Name	Description
device to mount	<p>This field identifies one of the following:</p> <ul style="list-style-type: none"> <li>■ The block device name for a local legacy UFS file system (for example, /dev/dsk/c8t1d0s7).</li> <li>■ The resource name for a remote file system (for example, myserver:/export/home). After you add an entry for a remote system resource, be sure that the following service is enabled.</li> </ul> <pre># svcs -a   grep nfs/client disabled      May_14   svc:/network/nfs/client:default # svcadm enable svc:/network/nfs/client:default</pre> <p>For more information about NFS, see <a href="#">Managing Network File Systems in Oracle Solaris 11.1</a>.</p> <ul style="list-style-type: none"> <li>■ The swap volume (for example, /dev/zvol/dsk/rpool/swap).</li> <li>■ A directory for a virtual file system.</li> </ul>
device to fsck	<p>The raw (character) device name that corresponds to the legacy UFS file system identified by the device to mount field (for example, /dev/rdisk/c8t1d0s7). This field determines the raw interface that is used by the fsck command. Use a dash (-) when there is no applicable device, such as for a read-only file system or a remote file system.</p>
mount point	<p>Identifies where to mount the legacy or remote file system (for example, /data).</p>
FS type	<p>Identifies the type of file system.</p>

TABLE 15-1 Field Descriptions for the `/etc/vfstab` File (Continued)

Field Name	Description
<code>fsck pass</code>	<p>The pass number used by the <code>fsck</code> command to decide whether to check a legacy UFS file system. When the field contains a dash (-), the file system is not checked. Currently, the <code>fsck pass</code> value in the <code>/etc/vfstab</code> file is ignored during the boot process.</p> <p>When the field contains a zero, legacy UFS file systems are not checked. When the field contains a value greater than zero, the UFS file system is always checked.</p> <p>All legacy UFS file systems with a value of 1 in this field are checked one at a time in the order they appear in the <code>vfstab</code> file. When the <code>fsck</code> command is run on multiple UFS file systems that have <code>fsck pass</code> values greater than 1 and the <code>preen</code> option (-o p) is used, the <code>fsck</code> command automatically checks the file systems on different disks in parallel to maximize efficiency. Otherwise, the value of the pass number does not have any effect.</p>
<code>mount at boot</code>	<p>Set to yes or no for whether the file system should be automatically mounted by the <code>mountall</code> command when the system is booted. Note that this field has nothing to do with <code>autofs</code>. This field should always be set to no for virtual file systems such as <code>/proc</code> and <code>/dev/fd</code>.</p>
<code>mount options</code>	<p>A list of comma-separated options (with no spaces) that are used for mounting the file system. Use a dash (-) to indicate no options. For more information, see <a href="#"><code>vfstab(4)</code></a>.</p>

**Note** – You must have an entry in each field in the `/etc/vfstab` file. If there is no value for a field, be sure to specify a dash (-). Otherwise, the system might not boot successfully. Similarly, white space should not be used as a field value.

## Prerequisites for Unmounting Oracle Solaris File Systems

The prerequisites for unmounting file systems include the following:

- You must be an administrator.
- You cannot unmount a file system that is busy. A file system is considered busy if a user is accessing a directory in the file system, if a program has a file open in that file system, or if the file system is being shared.

You can make a file system available for unmounting by doing the following:

- Changing to a directory in a different file system.
- Logging out of the system.

- Using the `fuser` command to list all processes that are accessing the file system and to stop them, if necessary. For more details, see [“How to Stop All Processes That Are Accessing a File System”](#) on page 306.

Notify users if you need to unmount a file system that they are using.

- Unsharing the file system.
  - For example:
 

```
zfs set share.nfs=off tank/fs1
```
  - Use the legacy unshare method. For information, see [unshare\(1M\)](#).

To verify that you unmounted a file system or a number of file systems, examine the output from the `mount` command:

```
$ mount | grep unmounted-file-system
```

## Creating and Mounting Oracle Solaris File Systems

This section provides examples of creating and mounting Oracle Solaris file systems.

### ▼ How to Create an ZFS File System

#### 1 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights”](#) in *Oracle Solaris 11.1 Administration: Security Services*.

#### 2 Create a ZFS storage pool.

The following example illustrates how to create a simple mirrored storage pool named `tank` and a ZFS file system named `tank` in one command. Assume that the whole disks `/dev/dsk/c1t0d0` and `/dev/dsk/c2t0d0` are available for use.

```
zpool create tank mirror c1t0d0 c2t0d0
```

#### 3 Create a ZFS file system.

```
zfs create tank/fs
```

The new ZFS file system, `tank/fs`, can use as much of the disk space as needed, and is automatically mounted at `/tank/fs`.

#### 4 Confirm that the file system is created.

```
zfs list -r tank
NAME USED AVAIL REFER MOUNTPOINT
tank 117K 268G 21K /tank
tank/fs 21K 268G 21K /tank/fs
```

## ▼ How to Create and Mount a Legacy UFS File System

**Before You Begin** Ensure that you have met the following prerequisites:

- The disk must be formatted and divided into slices.
- If you are recreating an existing legacy UFS file system, unmount it.
- You need to know the device name of the slice that will contain the file system.

For information on finding disks and disk slice numbers, see [Chapter 9, “Administering Disks \(Tasks\)”](#).

For information on formatting disks and dividing disks into slices, see [Chapter 7, “Managing Disks \(Overview\)”](#).

### 1 Become an administrator.

For more information, see “How to Use Your Assigned Administrative Rights” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Create a legacy UFS file system.

```
newfs [-N] [-b size] [-i bytes] /dev/rdisk/device-name
```

The system asks for confirmation.




---

**Caution** – Be sure you have specified the correct device name for the slice before performing this step. If you specify the wrong slice, you will erase its contents when the new file system is created. This error might cause the system to panic.

---

### 3 To verify the creation of the legacy UFS file system, check the new file system.

```
fsck /dev/rdisk/device-name
```

where *device-name* argument specifies the name of the disk device that contains the new file system.

The `fsck` command checks the consistency of the new file system, reports any problems, and prompts you before it repairs the problems. For more information on the `fsck` command, see [fsck\(1M\)](#).

### 4 Mount the legacy UFS file system.

```
mkdir /directory-name
mount /dev/dsk/device-name /directory-name
```

## Example 15–2 Creating and Mounting a Legacy UFS File System

The following example shows how to create and mount a UFS file system `/dev/rdisk/c0t1d0s0` on `/legacy`.

```

newfs /dev/rdisk/c0t1d0s0
newfs: construct a new file system /dev/rdisk/c0t1d0s0: (y/n)? y
/dev/rdisk/c0t1d0s0: 286722656 sectors in 46668 cylinders of 48 tracks, 128 sectors
 140001.3MB in 2917 cyl groups (16 c/g, 48.00MB/g, 5824 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
 32, 98464, 196896, 295328, 393760, 492192, 590624, 689056, 787488, 885920,
Initializing cylinder groups:
.....
super-block backups for last 10 cylinder groups at:
 285773216, 285871648, 285970080, 286068512, 286166944, 286265376, 286363808,
 286462240, 286560672, 286659104
fsck /dev/rdisk/c0t1d0s0
mkdir /legacy
mount /dev/dsk/c0t1d0s0 /legacy

```

### More Information After You Create a Legacy UFS File System ...

To mount the legacy UFS file system automatically at boot time, go to [“How to Add an Entry to the /etc/vfstab File”](#) on page 302.

## ▼ How to Create and Mount a TMPFS File System

### 1 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights”](#) in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Create the directory that you want to mount as the TMPFS file system, if necessary.

```
mkdir /mount-point
```

where *mount-point* is the directory on which the TMPFS file system is mounted.

### 3 Mount the TMPFS file system.

```
mount -F tmpfs [-o size=number] swap mount-point
```

*-o size=number* Specifies the size limit of the TMPFS file system in MB.

*mount-point* Specifies the directory on which the TMPFS file system is mounted.

To set up the system to automatically mount a TMPFS file system at boot time, see [Example 15-4](#).

### 4 Verify that the TMPFS file system has been created.

```
mount -v
```

### Example 15-3 Creating and Mounting a TMPFS File System

The following example shows how to create, mount, and limit the size of the TMPFS file system, /export/reports, to 50 MB.

```
mkdir /export/reports
chmod 777 /export/reports
mount -F tmpfs -o size=50m swap /export/reports
mount -v
```

### Example 15–4 Mounting a TMPFS File System at Boot Time

You can set up the system to automatically mount a TMPFS file system at boot time by adding an `/etc/vfstab` entry. The following example shows an entry in the `/etc/vfstab` file that mounts `/export/test` as a TMPFS file system at boot time. Because the `size=number` option is not specified, the size of the TMPFS file system on `/export/test` is limited only by the available system resources.

```
swap - /export/test tmpfs - yes -
```

## ▼ How to Create and Mount an LOFS File System

### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Create the directory you want to mount as an LOFS file system, if necessary.

```
mkdir loopback-directory
```

### 3 Grant the appropriate permissions and ownership on the newly created directory.

### 4 Create the mount point where you want to mount the LOFS file system, if necessary.

```
mkdir /mount-point
```

### 5 Mount the LOFS file system.

```
mount -F lofs loopback-directory /mount-point
```

`loopback-directory` Specifies the file system to be mounted on the loopback mount point.

`/mount-point` Specifies the directory on which to mount the LOFS file system.

### 6 Verify that the LOFS file system has been mounted.

```
mount -v
```

### Example 15–5 Creating and Mounting an LOFS File System

The following example shows how to create, mount, and test new software in the `/new/dist` directory as a loopback file system without actually having to install it.

```
mkdir /tmp/newroot
mount -F lofs /new/dist /tmp/newroot
chroot /tmp/newroot newcommand
```

### Example 15–6 Mounting an LOFS File System at Boot Time

You can set up the system to automatically mount an LOFS file system at boot time by adding an entry to the end of the `/etc/vfstab` file. The following example shows an entry in the `/etc/vfstab` file that mounts an LOFS file system for the root (`/`) file system on `/tmp/newroot`.

```
/ - /tmp/newroot lofs - yes -
```

Ensure that the loopback entries are the last entries in the `/etc/vfstab` file. Otherwise, if the `/etc/vfstab` entry for a loopback file system precedes the file systems to be included in it, the loopback file system cannot be mounted.

## ▼ How to Add an Entry to the `/etc/vfstab` File

Use this procedure to mount non-ZFS file systems at boot time unless legacy mount behavior is needed for some ZFS file systems. For more information about mounting ZFS file systems, see [Oracle Solaris 11.1 Administration: ZFS File Systems](#).

### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in [Oracle Solaris 11.1 Administration: Security Services](#).

### 2 Create a mount point for the file system to be mounted, if necessary.

```
mkdir /mount-point
```

There must be a mount point on the local system to mount a file system. A *mount point* is a directory to which the mounted file system is attached.

### 3 Edit the `/etc/vfstab` file and add an entry. Ensure that you do the following:

- a. Separate each field with white space (a space or a tab).
- b. Specify a dash (-) if a field has no contents.
- c. Save the changes.

---

**Note** – Because the root (`/`) file system is mounted read-only by the kernel during the boot process, only the `remount` option (and options that can be used in conjunction with `remount`) affect the root (`/`) entry in the `/etc/vfstab` file.

---

**Example 15-7** Adding an Entry to the `/etc/vfstab` File

The following example shows how to mount the disk slice `/dev/dsk/c0t3d0s7` as a legacy UFS file system to the mount point `/files1`. The raw character device `/dev/rdisk/c0t3d0s7` is specified as the device to `fsck`. The `fsck` pass value of 2 means that the file system will be checked, but not sequentially.

```
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
#
/dev/dsk/c0t3d0s7 /dev/rdisk/c0t3d0s7 /files1 ufs 2 yes -
```

The following example shows how to mount the `/export/man` directory from the system `pluto` as an NFS file system on mount point `/usr/man`. Neither a device to `fsck` nor a `fsck` pass is specified because it's an NFS file system. In this example, mount options are `ro` (read-only) and `soft`.

```
#device device mount FS fsck mount mount
#to mount to fsck point type pass at boot options
pluto:/export/man - /usr/man nfs - yes ro,soft
```

After you add the remote system and resource to the `/etc/vfstab` file, be sure that the following service is started.

```
svcs -a | grep nfs/client
disabled May_14 svc:/network/nfs/client:default
svcadm enable svc:/network/nfs/client:default
```

Otherwise, the remote file system will not be mounted after the system is rebooted.

## ▼ How to Mount a File System (`/etc/vfstab` File)

### 1 Become an administrator.

For more information, see [“How to Use Your Assigned Administrative Rights” in Oracle Solaris 11.1 Administration: Security Services](#).

### 2 Mount a file system listed in the `/etc/vfstab` file.

```
mount /mount-point
```

where `/mount-point` specifies an entry in the mount point or device to mount field in the `/etc/vfstab` file. It is usually easier to specify the mount point.

**Example 15-8** Mounting a File System (`/etc/vfstab` File)

The following example shows how to mount the local `/legacy` file system that is listed in the `/etc/vfstab` file.

```
mount /legacy
```

### Example 15–9 Mounting All File Systems (/etc/vfstab File)

The following example shows the messages that are displayed when you use the `mountall` command and the file systems are already mounted.

```
mountall
mount: /tmp is already mounted or swap is busy
```

The following example shows how to mount all the local systems that are listed in the `/etc/vfstab` file.

```
mountall -l
```

The following example shows how to mount all available ZFS file systems.

```
zfs mount -a
```

The following example shows how to mount all the remote file systems that are listed in the `/etc/vfstab` file.

```
mountall -r
```

## ▼ How to Mount an NFS File System (mount Command)

### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Create a mount point for the file system to be mounted, if necessary.

```
mkdir /mount-point
```

There must be a mount point on the local system to mount a file system. A mount point is a directory to which the mounted file system is attached.

### 3 Ensure that the resource (file or directory) is available from a server.

To mount an NFS file system, the resource must be made available on the server by using the `share` command. For information on how to share resources, see “[About the NFS Service](#)” in *Oracle Solaris Administration: Network Services*.

### 4 Mount the NFS file system.

```
mount -F nfs [-o mount-options] server:/directory /mount-point
```

**Example 15–10** Mounting an NFS File System (mount Command)

The following example shows how to mount the `/export/packages` directory on `/mnt` from the server `pluto`.

```
mount -F nfs pluto:/export/packages /mnt
```

## ▼ x86: How to Mount a PCFS (DOS) File System From a Hard Disk (mount Command)

Use the following procedure to mount a PCFS (DOS) file system from a hard disk.

### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 Create a mount point for the file system to be mounted, if necessary.

```
mkdir /mount-point
```

There must be a mount point on the local system to mount a file system. A *mount point* is a directory to which the mounted file system is attached.

### 3 Mount the PCFS file system.

```
mount -F pcfs [-o rw | ro] /dev/dsk/device-name:logical-drive /mount-point
```

`-o rw | ro` Specifies that you can mount a PCFS file system read/write (`rw`) or read-only (`ro`). If you do not specify this option, the default is `rw`.

`/dev/dsk/device-name` Specifies the device name of the whole disk (for example, `/dev/dsk/c0t0d0p0`).

`logical-drive` Specifies either the DOS logical drive letter (`c` through `z`) or a drive number (`1` through `24`). Drive `c` is equivalent to drive `1` and represents the primary DOS slice on the drive. All other letters or numbers represent DOS logical drives within the extended DOS slice.

`/mount-point` Specifies the directory on which to mount the file system.

Note that the *device-name* and *logical-drive* must be separated by a colon.

**Example 15–11** x86: Mounting a PCFS (DOS) File System From a Hard Disk (mount Command)

The following example shows how to mount the logical drive in the primary DOS slice on the `/pcfs/c` directory.

```
mount -F pcfs /dev/dsk/c0t0d0p0:c /pcfs/c
```

The following example shows how to mount read-only the first logical drive in the extended DOS slice on the /mnt directory.

```
mount -F pcfs -o ro /dev/dsk/c0t0d0p0:2 /mnt
```

## ▼ How to Stop All Processes That Are Accessing a File System

### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

### 2 List all the processes that are accessing the file system so that you know which processes you are going to stop.

```
fuser -c [-u] /mount-point
```

-c                    Reports on files that are mount points for file systems and any files within those mounted file systems.

-u                    Displays the user login name for each process ID.

/mount-point        Specifies the name of the file system for which you want to stop processes.

### 3 Stop all processes that are accessing the file system.

```
fuser -c -k /mount-point
```

A SIGKILL is sent to each process that is using the file system.

---

**Note** – You should not stop a user's processes without first warning the user.

---

### 4 Verify that no processes are accessing the file system.

```
fuser -c /mount-point
```

## Example 15–12 Stopping All Processes That Are Accessing a File System

The following example shows how to stop process 4006c that is using the /export/home file system.

```
fuser -c /export/home
/export/home: 4006c
fuser -c -k /export/home
/export/home: 4006c
```

```
fuser -c /export/home
/export/home:
```

## ▼ How to Unmount a File System

Use the following procedure to unmount a file system.

- 1 **Ensure that you have met the prerequisites listed in “Prerequisites for Unmounting Oracle Solaris File Systems” on page 297.**

- 2 **Unmount the file system.**

```
umount /mount-point
```

where */mount-point* is the name of the file system that you want to unmount.

This can be one of the following:

- The directory name where the file system is mounted
- The device name path of the file system
- The resource for an NFS file system
- The loopback directory for an LOFS file system

### Example 15–13 Unmounting a File System

The following example shows how to unmount a legacy UFS file system:

```
umount /legacy
```

The following example shows how to forcibly unmount the UFS */legacy* file system:

```
umount -f /legacy
#
```

The following example shows to unmount all ZFS file systems:

```
zfs umount -a
```

All file systems are unmounted, except for those file systems that are busy.



## Configuring Additional Swap Space (Tasks)

---

This chapter provides guidelines and step-by-step instructions for configuring additional swap space for a ZFS root file system after the Oracle Solaris OS is installed.

This is a list of the information in this chapter:

- “About Swap Space” on page 309
- “How Do I Know If I Need More Swap Space?” on page 312
- “How Swap Space Is Allocated” on page 313
- “Planning for Swap Space” on page 314
- “Monitoring Swap Resources” on page 315
- “Adding or Changing Swap Space in an Oracle Solaris ZFS Root Environment” on page 316

### About Swap Space

You should understand the features of the swap mechanism in Oracle Solaris to determine the following:

- Swap space requirements
- The relationship between swap space and the TMPFS file system
- How to recover from error messages related to swap space

### Swap Space and Virtual Memory

Oracle Solaris OS software and application software can use some disk space for temporary storage rather than for file systems. The reserved area of the disk is called *swap* space. Swap space is used as virtual memory storage areas when the system does not have enough physical memory to handle current processes. In a ZFS root file system, the disk space reserved for swap is a ZFS volume.

The virtual memory system maps physical copies of files on disk to virtual addresses in memory. Physical memory pages that contain the data for these mappings can be backed by regular files in the file system, or by swap space. If the memory is backed by swap space it is referred to as *anonymous memory* because no identity is assigned to the disk space that is backing the memory.

The Oracle Solaris OS uses the concept of *virtual swap space*, a layer between anonymous memory pages and the physical storage (or disk-backed swap space) that actually back these pages. A system's virtual swap space is equal to the sum of all its physical (disk-backed) swap space plus a portion of the currently available physical memory.

Virtual swap space has these advantages:

- The need for large amounts of physical swap space is reduced because virtual swap space does not necessarily correspond to physical (disk) storage.
- A pseudo file system called SWAPFS provides addresses for anonymous memory pages. Because SWAPFS controls the allocation of memory pages, it has greater flexibility in deciding what happens to a page. For example, SWAPFS might change the page's requirements for disk-backed swap storage.

## Swap Space and the TMPFS File System

The TMPFS file system is activated automatically in the Oracle Solaris environment by an entry in the `/etc/vfstab` file. The TMPFS file system stores files and their associated information in memory (in the `/tmp` directory) rather than on disk, which speeds access to those files. This feature results in a major performance enhancement for applications such as compilers and DBMS products that use `/tmp` heavily.

The TMPFS file system allocates space in the `/tmp` directory from the system's swap resources. This feature means that as you use up space in the `/tmp` directory, you are also using up swap space. So, if your applications use the `/tmp` directory heavily and you do not monitor swap space usage, your system could run out of swap space.

Do use the following if you want to use TMPFS, but your swap resources are limited:

- Mount the TMPFS file system with the `size` option (`-o size`) to control how much swap resources TMPFS can use.
- Use your compiler's `TMPDIR` environment variable to point to another larger directory. Using your compiler's `TMPDIR` variable only controls whether the compiler is using the `/tmp` directory. This variable has no effect on other programs' use of the `/tmp` directory.

## Swap Space and Dump Device Configuration

A *dump device* is usually disk space that is reserved to store system crash dump information. When a system is installed, a ZFS swap volume and dump volume are created automatically. You can change a system's dump volume by using the `dumpadm` command. For more information, see [Chapter 1, “Managing System Crash Information \(Tasks\),” in \*Troubleshooting Typical Issues in Oracle Solaris 11.1\*](#).

In a ZFS root environment, swap and dump are configured as separate ZFS volumes. The advantages to this model are as follows:

- You don't have to partition a disk to include swap and dump areas.
- Swap and dump devices benefit from the underlying ZFS I/O pipeline architecture.
- You can set characteristics, such as compression, on swap and dump devices.
- You can reset swap and dump device sizes. For example:

```
zfs set volsize=2G rpool/dump
zfs get volsize rpool/dump
NAME PROPERTY VALUE SOURCE
rpool/dump volsize 2G -
```

Keep in mind that reallocating a large dump device is a time-consuming process.

- On an active system, you might consider adding a second swap volume to increase overall swap size.

For more information about using ZFS swap and dump devices, see [“Managing Your ZFS Swap and Dump Devices” in \*Oracle Solaris 11.1 Administration: ZFS File Systems\*](#).

## Swap Space and Dynamic Reconfiguration

A good practice is to allocate enough swap space to support a failing CPU or system board during dynamic reconfiguration. Otherwise, a CPU or system board failure might result in your host or domain rebooting with less memory.

Without having this additional swap space available, one or more of your applications might fail to start due to insufficient memory. This problem would require manual intervention either to add additional swap space or to reconfigure the memory usage of these applications.

If you have allocated additional swap space to handle a potential loss of memory on reboot, all of your intensive applications might start as usual. This means the system will be available to the users, perhaps possibly slower due to some additional swapping.

For more information, see your hardware dynamic reconfiguration guide.

## Configuring Swap Space in a SAN Environment

Review the following points to determine whether you might configure swap space on a network-connected disk, such as in a SAN environment:

- Diagnosing swap space issues on a locally-attached disk is easier than diagnosing swap space issues on a network-connected disk.
- The performance of swap space over a SAN should be comparable to swap space configured on a locally-attached disk.
- Adding more memory to a system with performance issues, after analyzing performance data, might resolve a swap over SAN performance problem better than moving the swap to a locally-attached disk.

## How Do I Know If I Need More Swap Space?

Use the `swap -l` command to determine if your system needs more swap space.

For example, the following `swap -l` output shows that this system's swap space is almost entirely consumed or at 100% allocation.

```
$ swap -l
swap -l
swapfile dev swaplo blocks free
/dev/zvol/dsk/rpool/swap 182,2 16 67108848 67102992
```

When a system's swap space is at 100% allocation, an application's memory pages become temporarily locked. Application errors might not occur, but system performance will likely suffer.

## Swap-Related Error Messages

These messages indicate that an application was trying to get more anonymous memory. However, no swap space was left to back it.

```
application is out of memory
```

```
malloc error 0
```

```
messages.1:Sep 21 20:52:11 mars genunix: [ID 470503 kern.warning]
WARNING: Sorry, no swap space to grow stack for pid 100295 (myprog)
```

## TMPFS-Related Error Messages

The following message is displayed if a page could not be allocated when a file was being written. This problem can occur when TMPFS tries to write more than it is allowed or if currently executed programs are using a lot of memory.

*directory:* File system full, swap space limit exceeded

The following message means that TMPFS ran out of physical memory while attempting to create a new file or directory:

*directory:* File system full, memory allocation failed

For information on recovering from the TMPFS-related error messages, see [tmpfs\(7FS\)](#).

## How Swap Space Is Allocated

Initially, swap space is allocated as part of the Oracle Solaris installation process. In a ZFS root environment, the default swap size is generally in the 512 MB to 2 GB range.

For general guidelines on allocating swap space, see “[Planning for Swap Space](#)” on page 314.

## Swap Areas and the `/etc/vfstab` File

After the system is installed, swap areas and swap files are listed in the `/etc/vfstab` file. They are activated by the `/sbin/swapadd` script when the system is booted.

An entry for a swap device in the `/etc/vfstab` file contains the following:

- The full path name of the swap volume path name on a system with a ZFS root file system
- File system type of the swap slice or swap file

The file system that contains a swap file must be mounted before the swap file is activated. So, in the `/etc/vfstab` file, ensure that the entry that mounts the file system comes before the entry that activates the swap file.

# Planning for Swap Space

The most important factors in determining swap space size are the requirements of the system's software applications. For example, large applications such as computer-aided design simulators, database management products, transaction monitors, and geologic analysis systems can consume as much as 200–1000 MB of swap space.

Consult your application vendors for swap space requirements for their applications.

If you are unable to determine swap space requirements from your application vendors, use the following general guidelines based on your system type to allocate swap space.

---

**Note** – Crash dump content is compressed so the dump device does not have to be the same size as physical memory. By default, the dump content value is set to kernel pages. However, if the dump content value is set to dump all memory pages, then consider increasing the dump size to half the size of physical memory or more.

---

**TABLE 16-1** Swap and Dump Volume Sizes for ZFS File Systems

System Type	Swap Volume Size	Dump Volume Size
System with about 4 GB of physical memory	1 GB	2 GB
Mid-range server with about 8 GB of physical memory	2 GB	4 GB
High-end server with about 16 to 128 GB of physical memory	4 GB	8-64 GB
High-end server with more than 128 GB of physical memory	1/4 of physical memory size	1/2 of physical memory size

---

**Note** – A busy system with many active ZFS file systems might use 1/2 to 3/4 the size of physical memory for the size of the dump device.

On a system with a ZFS root file system, if you attempt to designate a dump device that is too small to hold a system crash system with the `dumpadm -d` command, you will see a message similar to the following:

```
dumpadm: dump device /dev/zvol/dsk/rpool/dump is too small to hold a
system dump dump size 43467329536 bytes, device size 42949672960 bytes
```

---

## Allocating Swap Space for ZFS-Based Systems

During an initial installation of a ZFS root file system, a swap area is automatically created on a ZFS volume in the ZFS root pool.

In a ZFS root pool, swap devices are not pre-allocated to fixed-size slices, so it is fairly easy to modify the swap size later.

After you assess the swap requirements of your applications, you can use the default swap size or adjust the swap volume size during an initial installation or after the installation, if necessary.

During an initial installation, the default dump volume size is calculated by the kernel based on `dumppadm` information and the size of physical memory.

In a ZFS environment, file systems consume space from the pool so the `/var/crash` directory consumes what it needs depending on how many crash dumps are saved.

## Monitoring Swap Resources

The `/usr/sbin/swap` command is used to manage swap areas. Two options, `-l` and `-s`, display information about swap resources.

Use the `swap -l` command to identify a system's swap areas. Activated swap devices or files are listed under the `swapfile` column. For example:

```
swap -l
swapfile dev swaplo blocks free
/dev/dsk/c0t0d0s1 136,1 16 1638608 1600528
```

On a system with a ZFS root file system, the `swap -l` command identifies similar output except that it identifies the ZFS volume path name. For example:

```
swap -l
swapfile dev swaplo blocks free
/dev/zvol/dsk/rpool/swap 256,1 16 1058800 1058800
```

Use the `swap -s` command to monitor swap resources.

```
swap -s
total: 57416k bytes allocated + 10480k reserved = 67896k used,
833128k available
```

The `used` value plus the `available` value equals the total swap space on the system, which includes a portion of physical memory and swap devices (or files).

You can use the amount of available and used swap space (in the `swap -s` output) as a way to monitor swap space usage over time. If a system's performance is good, use `swap -s` to determine how much swap space is available. When the performance of a system slows down,

check the amount of available swap space to determine if it has decreased. Then you can identify what changes to the system might have caused swap space usage to increase.

When using this command, keep in mind that the amount of physical memory available for swap usage changes dynamically as the kernel and user processes lock down and release physical memory.

---

**Note** – The `swap -l` command displays swap space in 512-byte blocks. The `swap -s` command displays swap space in 1024-byte blocks. If you add up the blocks from `swap -l` and convert them to KB, the result is less than `used + available` (in the `swap -s` output). The reason is that `swap -l` does not include physical memory in its calculation of swap space.

---

The output from the `swap -s` command is summarized in the following table.

TABLE 16-2 Output of the `swap -s` Command

Keyword	Description
bytes allocated	The total amount of swap space in 1024-byte blocks that is currently allocated as backing store (disk-backed swap space).
reserved	The total amount of swap space in 1024-byte blocks that is not currently allocated, but claimed by memory for possible future use.
used	The total amount of swap space in 1024-byte blocks that is either allocated or reserved.
available	The total amount of swap space in 1024-byte blocks that is currently available for future reservation and allocation.

## Adding or Changing Swap Space in an Oracle Solaris ZFS Root Environment

The following section describes how to add or change swap space in a ZFS root environment. See the previous sections to determine if your system or applications need additional swap space.

### ▼ How to Add Swap Space in an Oracle Solaris ZFS Root Environment

#### 1 Become an administrator.

For more information, see “[How to Use Your Assigned Administrative Rights](#)” in *Oracle Solaris 11.1 Administration: Security Services*.

**2 Identify the current swap volume.**

A swap volume cannot be removed if it is in use. You can tell if the current swap volume is in use by comparing the blocks identified in the blocks column and blocks identified in the free column. If the blocks in the two columns are equal, the swap area is not busy. For example:

```
swap -l
swapfile dev swaplo blocks free
/dev/zvol/dsk/rpool/swap 256,1 16 1058800 1058800
```

**3 Select one of the following to resize the swap volume.**

- a. If the current swap area is not in use, you can resize the size of the current swap volume, but you must reboot the system to see the increased swap space.**

For example:

```
zfs get volsize rpool/swap
NAME PROPERTY VALUE SOURCE
rpool/swap volsize 517M -
zfs set volsize=2g rpool/swap
zfs get volsize rpool/swap
NAME PROPERTY VALUE SOURCE
rpool/swap volsize 2G -
init 6
```

- b. If the system cannot be rebooted, add another swap volume to increase your total swap space.**

For example:

```
zfs create -V 2G rpool/swap2
```

Then, activate the second swap volume.

```
swap -a /dev/zvol/dsk/rpool/swap2
swap -l
swapfile dev swaplo blocks free
/dev/zvol/dsk/rpool/swap 256,1 16 1058800 1058800
/dev/zvol/dsk/rpool/swap2 256,3 16 4194288 4194288
```

**4 If necessary, add an entry for the second swap volume in the `/etc/vfstab` file.**

For example:

```
/dev/zvol/dsk/rpool/swap2 - - swap - no -
```



## Copying Files and File Systems (Tasks)

---

This chapter describes how to copy files and file systems to disk, tape, and diskettes by using various backup commands.

This is a list of the step-by-step instructions in this chapter.

- “How to Copy Directories Between File Systems (`cpio`)” on page 322
- “How to Copy Files to a Tape (`tar`)” on page 323
- “How to List the Files on a Tape (`tar`)” on page 324
- “How to Retrieve Files From a Tape (`tar`)” on page 325
- “Copying Files to a Tape With the `pax` Command” on page 326
- “How to Copy All Files in a Directory to a Tape (`cpio`)” on page 327
- “How to List the Files on a Tape (`cpio`)” on page 328
- “How to Retrieve All Files From a Tape (`cpio`)” on page 328
- “How to Retrieve Specific Files From a Tape (`cpio`)” on page 329
- “How to Copy Files to a Remote Tape Device (`tar` and `dd`)” on page 330
- “How to Extract Files From a Remote Tape Device” on page 331

### Commands for Copying File Systems

When you want to copy or move individual files, portions of file systems, or complete file systems, you can use the procedures described in this chapter.

The following table describes various backup and restore commands that are available in the Oracle Solaris release. For enterprise environments, consider using an enterprise-level backup product. Information about enterprise-level backup products is available on the Oracle Technical Network.

**TABLE 17-1** Summary of Various Backup Commands

Command Name	Aware of File System Boundaries?	Supports Multiple Volume Backups?	Physical or Logical Copy?
volcopy	Yes	Yes	Physical
tar	No	No	Logical
cpio	No	Yes	Logical
pax	Yes	Yes	Logical
dd	Yes	No	Physical
zfs send and zfs receive	Yes	N/A	Logical
zfs snapshot	Yes	N/A	Logical

The following table describes the advantages and disadvantages of some of these commands.

**TABLE 17-2** Advantages and Disadvantages of tar, pax, and cpio Commands

Command	Function	Advantages	Disadvantages
tar	Use to copy files and directory subtrees to a single tape.	<ul style="list-style-type: none"> <li>■ Available on most UNIX operating systems</li> <li>■ Public domain versions are readily available</li> </ul>	<ul style="list-style-type: none"> <li>■ Is not aware of file system boundaries</li> <li>■ Length of full path name cannot exceed 255 characters</li> <li>■ Cannot be used to create multiple tape volumes</li> </ul>
pax	Use to copy files, special files, or file systems that require multiple tape volumes. Or, use when you want to copy files to and from POSIX-compliant systems.	<ul style="list-style-type: none"> <li>■ Better portability than the tar or cpio commands for POSIX-compliant systems</li> <li>■ Multiple vendor support</li> </ul>	Same disadvantages as the tar command, except that the pax command can create multiple tape volumes.

TABLE 17-2 Advantages and Disadvantages of tar, pax, and cpio Commands (Continued)

Command	Function	Advantages	Disadvantages
cpio	Use to copy files, special files, or file systems that require multiple tape volumes. Or, use when you want to copy files from systems running current Oracle Solaris releases systems to systems running older Solaris releases.	<ul style="list-style-type: none"> <li>■ Packs data onto tape more efficiently than the tar command</li> <li>■ Skips over any bad spots in a tape when restoring</li> <li>■ Provides options for writing files with different header formats, such as ( tar, ustar, crc, odc, bar), for portability between different system types</li> <li>■ Creates multiple tape volumes</li> </ul>	The command syntax is more difficult than the tar or pax commands.

The following sections describes step-by-step instructions and examples of how to use these commands.

## Copying Directories Between File Systems (cpio Command)

You can use the `cpio` (copy in and out) command to copy individual files, groups of files, or complete file systems. This section describes how to use the `cpio` command to copy complete file systems.

The `cpio` command is an archiving program that copies a list of files into a single, large output file. This command inserts headers between the individual files to facilitate recovery. You can use the `cpio` command to copy complete file systems to another slice, another system, or to a media device, such as a tape or USB diskette.

Because the `cpio` command recognizes end-of-media and prompts you to insert another volume, it is the most effective command to use to create archives that require multiple tapes or USB diskettes.

With the `cpio` command, you frequently use the `ls` and `find` commands to list and select the files you want to copy, and then to pipe the output to the `cpio` command.

## ▼ How to Copy Directories Between File Systems (cpio)

1 Become an administrator.

2 Change to the appropriate directory.

```
cd filesystem1
```

3 Copy the directory tree from *filesystem1* to *filesystem2* by using a combination of the `find` and `cpio` commands.

```
find . -print -depth | cpio -pdm filesystem2
```

`.` Starts in the current working directory.

`-print` Prints the file names.

`-depth` Descends the directory hierarchy and prints file names from the bottom up.

`-p` Creates a list of files.

`-d` Creates directories as needed.

`-m` Sets the correct modification times on directories.

For more information, see [cpio\(1\)](#).

The files from the directory name you specify are copied. The symbolic links are preserved.

You might also specify the `-u` option. This option forces an unconditional copy. Otherwise, older files do not replace newer files. This option might be useful if you want an exact copy of a directory, and some of the files being copied might already exist in the target directory.

4 Verify that the copy was successful by displaying the contents of the destination directory.

```
cd filesystem2
ls
```

5 If appropriate, remove the source directory.

```
rm -rf filesystem1
```

### Example 17-1 Copying Directories Between File Systems (cpio)

```
cd /data1
find . -print -depth | cpio -pdm /data2
19013 blocks
cd /data2
ls
rm -rf /data1
```

## Copying Files and File Systems to Tape

You can use the `tar`, `pax`, and `cpio` commands to copy files and file systems to tape. The command that you choose depends on how much flexibility and precision you require for the copy. Because all three commands use the raw device, you do not need to format or make a file system on tapes before you use them.

The tape drive and device name that you use depend on the hardware configuration for each system. For more information about tape device names, see [“Choosing Which Media to Use”](#) on page 333.

## Copying Files to Tape (tar Command)

Here is information that you should know before you copy files to tape with the `tar` command:

- Copying files to a tape with the `-c` option to the `tar` command destroys any files already on the tape at or beyond the current tape position.
- You can use file name substitution wildcards (`?` and `*`) as part of the file names that you specify when copying files. For example, to copy all documents with a `.doc` suffix, type `*.doc` as the file name argument.
- You cannot use file name substitution wildcards when you extract files from a `tar` archive.

### ▼ How to Copy Files to a Tape (tar)

- 1 **Change to the directory that contains the files you want to copy.**
- 2 **Insert a write-enabled tape into the tape drive.**
- 3 **Copy the files to tape.**

```
$ tar cvf /dev/rmt/n filenames
```

<code>c</code>	Indicates that you want to create an archive.
<code>v</code>	Displays the name of each file as it is archived.
<code>f /dev/rmt/n</code>	Indicates that the archive should be written to the specified device or file.
<code>filenames</code>	Indicates the files and directories that you want to copy. Separate multiple files with spaces.

The file names that you specify are copied to the tape, overwriting any existing files on the tape.

- 4 **Remove the tape from the drive. Write the names of the files on the tape label.**

**5 Verify that the files you copied are on the tape.**

```
$ tar tvf /dev/rmt/n
```

For more information on listing files on a tar tape, see [“How to List the Files on a Tape \(tar\)” on page 324](#).

**Example 17-2 Copying Files to a Tape (tar)**

The following example shows how to copy three files to the tape in tape drive 0.

```
$ cd /export/home/kryten
$ ls reports
reportA reportB reportC
$ tar cvf /dev/rmt/0 reports
a reports/ 0 tape blocks
a reports/reportA 59 tape blocks
a reports/reportB 61 tape blocks
a reports/reportC 63 tape blocks
$ tar tvf /dev/rmt/0
```

**▼ How to List the Files on a Tape (tar)****1 Insert a tape into the tape drive.****2 Display the tape contents.**

```
$ tar tvf /dev/rmt/n
```

t Lists the table of contents for the files on the tape.

v Used with the t option, and provides detailed information about the files on the tape.

f /dev/rmt/n Indicates the tape device.

**Example 17-3 Listing the Files on a Tape (tar)**

The following example shows a listing of files on the tape in drive 0.

```
$ tar tvf /dev/rmt/0
drwxr-xr-x 0/0 0 Jul 14 13:50 2010 reports/
-r--r--r-- 0/0 206663 Jul 14 13:50 2010 reports/reportC
-r--r--r-- 0/0 206663 Jul 14 13:50 2010 reports/reportB
-r--r--r-- 0/0 206663 Jul 14 13:50 2010 reports/reportA
```

## ▼ How to Retrieve Files From a Tape (tar)

- 1 Change to the directory where you want to put the files.
- 2 Insert the tape into the tape drive.
- 3 Retrieve the files from the tape.

```
$ tar xvf /dev/rmt/n [filenames]
```

x Indicates that the files should be extracted from the specified archive file. All files on the tape in the specified drive are copied to the current directory.

v Displays the name of each file as it is retrieved.

f /dev/rmt/n Indicates the tape device that contains the archive.

*filenames* Specifies a file to retrieve. Separate multiple files with spaces.

For more information, see the [tar\(1\)](#) man page.

- 4 Verify that the files have been copied.

```
$ ls -l
```

### Example 17-4 Retrieving Files on a Tape (tar)

The following example shows how to retrieve all the files from the tape in drive 0.

```
$ cd /var/tmp
$ tar xvf /dev/rmt/0
x reports/, 0 bytes, 0 tape blocks
x reports/reportA, 0 bytes, 0 tape blocks
x reports/reportB, 0 bytes, 0 tape blocks
x reports/reportC, 0 bytes, 0 tape blocks
x reports/reportD, 0 bytes, 0 tape blocks
$ ls -l
```

**Troubleshooting** The names of the files extracted from the tape must exactly match the names of the files that are stored on the archive. If you have any doubts about the names or paths of the files, first list the files on the tape. For instructions on listing the files on the tape, see [“How to List the Files on a Tape \(tar\)” on page 324](#).

# Copying Files to a Tape With the pax Command

## ▼ How to Copy Files to a Tape (pax)

1 Change to the directory that contains the files you want to copy.

2 Insert a write-enabled tape into the tape drive.

3 Copy the files to tape.

```
$ pax -w -f /dev/rmt/n filenames
```

-w Enables the write mode.

-f /dev/rmt/n Identifies the tape drive.

*filenames* Indicates the files and directories that you want to copy. Separate multiple files with spaces.

For more information, see the [pax\(1\)](#) man page.

4 Verify that the files have been copied to tape.

```
$ pax -f /dev/rmt/n
```

5 Remove the tape from the drive. Write the names of the files on the tape label.

### Example 17-5 Copying Files to a Tape (pax)

The following example shows how to use the pax command to copy all the files in the current directory.

```
$ pax -w -f /dev/rmt/0 .
$ pax -f /dev/rmt/0
filea fileb filec
```

# Copying Files to Tape With the `cpio` Command

## ▼ How to Copy All Files in a Directory to a Tape (`cpio`)

- 1 Change to the directory that contains the files you want to copy.
- 2 Insert a write-enabled tape into the tape drive.
- 3 Copy the files to tape.

```
$ ls | cpio -oc > /dev/rmt/n
```

`ls` Provides the `cpio` command with a list of file names.

`cpio -oc` Specifies that the `cpio` command should operate in copy-out mode (`-o`) and write header information in ASCII character format (`-c`). These options ensure portability to other vendors' systems.

`> /dev/rmt/n` Specifies the output file.

All files in the directory are copied to the tape in the drive you specify, overwriting any existing files on the tape. The total number of blocks that are copied is shown.

- 4 Verify that the files have been copied to tape.

```
$ cpio -civt < /dev/rmt/n
```

`-c` Specifies that the `cpio` command should read files in ASCII character format.

`-i` Specifies that the `cpio` command should operate in copy-in mode, even though the command is only listing files at this point.

`-v` Displays the output in a format that is similar to the output from the `ls -l` command.

`-t` Lists the table of contents for the files on the tape in the tape drive that you specify.

`< /dev/rmt/n` Specifies the input file of an existing `cpio` archive.

- 5 Remove the tape from the drive. Write the names of the files on the tape label.

### Example 17-6 Copying All Files in a Directory to a Tape (`cpio`)

The following example shows how to copy all of the files in the `/export/home/kryten` directory to the tape in tape drive 0.

```

$ cd /export/home/kryten
$ ls | cpio -oc > /dev/rmt/0
1280 blocks
$ cpio -civt < /dev/rmt/0
-r--r--r-- 1 kryten staff 206663 Jul 14 13:52 2010, filea
-r--r--r-- 1 kryten staff 206663 Jul 14 13:52 2010, fileb
-r--r--r-- 1 kryten staff 206663 Jul 14 13:52 2010, filec
drwxr-xr-x 2 kryten staff 0 Jul 14 13:52 2010, letters
drwxr-xr-x 2 kryten staff 0 Jul 14 13:52 2010, reports
1280 blocks

```

## ▼ How to List the Files on a Tape (`cpio`)

---

**Note** – Listing the table of contents on a tape takes a long time because the `cpio` command must process the entire archive.

---

- 1 Insert an archive tape into the tape drive.
- 2 List the files on the tape.

```
$ cpio -civt < /dev/rmt/n
```

### Example 17-7 Listing the Files on a Tape (`cpio`)

The following example shows how to list the files on the tape in drive 0.

```

$ cpio -civt < /dev/rmt/0
-r--r--r-- 1 kryten staff 206663 Jul 14 13:52 2010, filea
-r--r--r-- 1 kryten staff 206663 Jul 14 13:52 2010, fileb
-r--r--r-- 1 kryten staff 206663 Jul 14 13:52 2010, filec
drwxr-xr-x 2 kryten staff 0 Jul 14 13:52 2010, letters
drwxr-xr-x 2 kryten staff 0 Jul 14 13:52 2010, reports
1280 blocks

```

## ▼ How to Retrieve All Files From a Tape (`cpio`)

If the archive was created using relative path names, the input files are built as a directory within the current directory when you retrieve the files. If, however, the archive was created with absolute path names, the same absolute paths are used to recreate the file on your system.




---

**Caution** – The use of absolute path names can be dangerous because you might overwrite existing files on your system.

---

- 1 Change to the directory where you want to put the files.

**2 Insert the tape into the tape drive.****3 Extract all files from the tape.**

```
$ cpio -icvd < /dev/rmt/n
```

-i Extracts files from standard input.

-c Specifies that the `cpio` command should read files in ASCII character format.

-v Displays the files as they are retrieved in a format that is similar to the output from the `ls` command.

-d Creates directories as needed.

< /dev/rmt/n Specifies the output file.

**4 Verify that the files were copied.**

```
$ ls -l
```

**Example 17-8 Retrieving All Files From a Tape (`cpio`)**

The following example shows how to retrieve all files from the tape in drive 0.

```
$ cd /var/tmp
cpio -icvd < /dev/rmt/0
answers
sc.directives
tests
8 blocks
$ ls -l
```

**▼ How to Retrieve Specific Files From a Tape (`cpio`)****1 Change to the directory where you want to put the files.****2 Insert the tape into the tape drive.****3 Retrieve a subset of files from the tape.**

```
$ cpio -icv "*file" < /dev/rmt/n
```

-i Extracts files from standard input.

-c Specifies that the `cpio` command should read headers in ASCII character format.

- v                    Displays the files as they are retrieved in a format that is similar to the output from the `ls` command.
- "\*file"             Specifies that all files that match the pattern are copied to the current directory. You can specify multiple patterns, but each pattern must be enclosed in double quotation marks.
- < /dev/rmt/n        Specifies the input file.

For more information, see the [cpio\(1\)](#) man page.

#### 4 Verify that the files were copied.

```
$ ls -l
```

### Example 17–9 Retrieving Specific Files From a Tape (cpio)

The following example shows how to retrieve all files with the `chapter` suffix from the tape in drive 0.

```
$ cd /home/smith/Book
$ cpio -icv "**chapter" < /dev/rmt/0
Boot.chapter
Directory.chapter
Install.chapter
Intro.chapter
31 blocks
$ ls -l
```

## Copying Files to a Remote Tape Device

### ▼ How to Copy Files to a Remote Tape Device (tar and dd)

- 1 Configure `ssh` on the remote system so that you can access the tape drive. See [“Configuring Secure Shell \(Tasks\)”](#) in *Oracle Solaris 11.1 Administration: Security Services*.
- 2 Change to the directory where you want to put the files.
- 3 Insert the tape into the tape drive.
- 4 Copy the files to a remote tape drive.

```
$ tar cvf - filenames | ssh remote-host dd of=/dev/rmt/n obs=block-size
```

<code>tar cf</code>	Creates a tape archive, lists the files as they are archived, and specifies the tape device.
<code>v</code>	Provides additional information about the tar file entries.
<code>-</code> (Hyphen)	Represents a placeholder for the tape device.
<i>filenames</i>	Identifies the files to be copied. Separate multiple files with spaces.
<code>ssh   remote-host</code>	Pipes the tar command's output to a remote system.
<code>dd of= /dev/rmt/n</code>	Represents the output device.
<code>obs=block-size</code>	Represents the blocking factor.

**5 Remove the tape from the drive. Write the names of the files on the tape label.**

**Example 17-10 Copying Files to a Remote Tape Drive (tar and dd)**

```
tar cvf - * | ssh mercury dd of=/dev/rmt/0 obs=126b
password:
a answers/ 0 tape blocks
a answers/test129 1 tape blocks
a sc.directives/ 0 tape blocks
a sc.directives/sc.190089 1 tape blocks
a tests/ 0 tape blocks
a tests/test131 1 tape blocks
6+9 records in
0+1 records out
```

## ▼ How to Extract Files From a Remote Tape Device

**1 Insert the tape into the tape drive.**

**2 Change to a temporary directory.**

```
$ cd /var/tmp
```

**3 Extract the files from a remote tape device.**

```
$ ssh remote-host dd if=/dev/rmt/n | tar xvBpf -
```

`ssh remote-host` Indicates a secure shell that is started to extract the files from the tape device by using the dd command.

`dd if=/dev/rmt/n` Indicates the input device.

`| tar xvBpf -` Pipes the output of the dd command to the tar command, which is used to restore the files.

**4 Verify that the files have been extracted.**

```
$ ls -l
```

**Example 17-11** Extracting Files From a Remote Tape Drive

```
$ cd /var/tmp
$ ssh mercury dd if=/dev/rmt/0 | tar xvBpf -
password:
x answers/, 0 bytes, 0 tape blocks
x answers/test129, 48 bytes, 1 tape blocks
20+0 records in
20+0 records out
x sc.directives/, 0 bytes, 0 tape blocks
x sc.directives/sc.190089, 77 bytes, 1 tape blocks
x tests/, 0 bytes, 0 tape blocks
x tests/test131, 84 bytes, 1 tape blocks
$ ls -l
```

## Managing Tape Drives (Tasks)

---

This chapter describes how to manage tape drives in the Oracle Solaris OS.

This is a list of the information in this chapter:

- “Choosing Which Media to Use” on page 333
- “Backup Device Names” on page 334
- “Displaying Tape Drive Status” on page 336
- “Handling Magnetic Tape Cartridges” on page 337
- “Guidelines for Drive Maintenance and Media Handling” on page 337

This is a list of the step-by-step instructions in this chapter.

- “How to Display Tape Drive Status” on page 336
- “Retensioning a Magnetic Tape Cartridge” on page 337
- “Rewinding a Magnetic Tape Cartridge” on page 337

### Choosing Which Media to Use

You typically back up Oracle Solaris systems by using the following tape media:

- 1/2-inch reel tape
- 1/4-inch streaming cartridge tape
- 8-mm cartridge tape
- 4-mm cartridge tape (DAT)

The media that you choose depends on the availability of the equipment that supports it and of the media (usually tape) that you use to store the files. Although you must do the backup from a local system, you can write the files to a remote device.

The following table shows typical tape devices that are used for backing up file systems. The storage capacity for each device depends on the type of drive and the data being written to the tape.

TABLE 18-1 Media Storage Capacities

Backup Media	Storage Capacity
1/2-inch reel tape	140 MB (6250 bpi)
2.5-GB 1/4-inch cartridge (QIC) tape	2.5 GB
DDS3 4-mm cartridge tape (DAT)	12–24 GB
14-GB 8-mm cartridge tape	14 GB
DLT 7000 1/2-inch cartridge tape	35–70 GB

## Backup Device Names

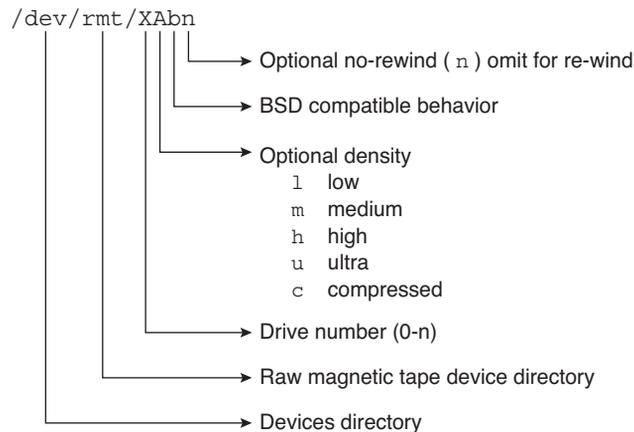
You specify a tape or device to use for backup by supplying a logical device name. This name points to the subdirectory that contains the “raw” device file and includes the logical unit number of the drive. Tape drive naming conventions use a logical, not a physical, device name. The following table shows this naming convention.

TABLE 18-2 Basic Device Names for Backup Devices

Device Type	Name
Tape	<code>/dev/rmt/n</code>

In general, you specify a tape device as shown in the following figure.

FIGURE 18-1 Tape Drive Device Names



If you don't specify the density, a tape drive typically writes at its “preferred” density. The preferred density usually means the highest density the tape drive supports. Most SCSI drives can automatically detect the density or format on the tape and read it accordingly. To determine the different densities that are supported for a drive, look at the `/dev/rmt` subdirectory. This subdirectory includes the set of tape device files that support different output densities for each tape.

Also, a SCSI controller can have a maximum of seven SCSI tape drives.

## Specifying the Rewind Option for a Tape Drive

Normally, you specify a tape drive by its logical unit number, which can run from 0 to *n*. The following table describes how to specify tape device names with a rewind or a no-rewind option.

TABLE 18-3 Specifying Rewind or No-Rewind for a Tape Drive

Drive and Rewind Value	Use This Option
First drive, rewind	<code>/dev/rmt/0</code>
First drive, no rewind	<code>/dev/rmt/0n</code>
Second drive, rewind	<code>/dev/rmt/1</code>
Second drive, no rewind	<code>/dev/rmt/1n</code>

## Specifying Different Densities for a Tape Drive

By default, the drive writes at its “preferred” density, which is usually the highest density the tape drive supports. If you do not specify a tape device, the command writes to drive number 0 at the default density the device supports.

To transport a tape to a system whose tape drive supports only a certain density, specify a device name that writes at the desired density. The following table describes how to specify different densities for a tape drive.

TABLE 18-4 Specifying Different Densities for a Tape Drive

Drive, Density, and Rewind Value	Use This Option
First drive, low density, rewind	<code>/dev/rmt/0l</code>
First drive, low density, no rewind	<code>/dev/rmt/0ln</code>
Second drive, medium density, rewind	<code>/dev/rmt/1m</code>
Second drive, medium density, no rewind	<code>/dev/rmt/1mn</code>

The additional density values are shown in “Backup Device Names” on page 334.

## Displaying Tape Drive Status

You can use the `status` option with the `mt` command to get status information about tape drives. The `mt` command reports information about any tape drives that are described in the `/kernel/drv/st.conf` file.

### ▼ How to Display Tape Drive Status

- 1 Load a tape into the drive you want information about.
- 2 Display the tape drive status.  

```
mt -f /dev/rmt/n status
```
- 3 Repeat steps 1–2, substituting tape drive numbers 0, 1, 2, 3, and so on to display information about all available tape drives.

#### Example 18-1 Displaying Tape Drive Status

The following example shows the status for a QIC-150 tape drive (`/dev/rmt/0`):

```
$ mt -f /dev/rmt/0 status
Archive QIC-150 tape drive:
 sense key(0x0)= No Additional Sense residual= 0 retries= 0
 file no= 0 block no= 0
```

The following example shows the status for an Exabyte tape drive (`/dev/rmt/1`):

```
$ mt -f /dev/rmt/1 status
Exabyte EXB-8200 8mm tape drive:
sense key(0x0)= NO Additional Sense residual= 0 retries= 0
file no= 0 block no= 0
```

The following example shows a quick way to poll a system and locate all of its tape drives:

```
$ for drive in 0 1 2 3 4 5 6 7
> do
> mt -f /dev/rmt/$drive status
> done
Archive QIC-150 tape drive:
 sense key(0x0)= No Additional Sense residual= 0 retries= 0
 file no= 0 block no= 0
/dev/rmt/1: No such file or directory
/dev/rmt/2: No such file or directory
/dev/rmt/3: No such file or directory
```

```
/dev/rmt/4: No such file or directory
/dev/rmt/5: No such file or directory
/dev/rmt/6: No such file or directory
/dev/rmt/7: No such file or directory
$
```

## Handling Magnetic Tape Cartridges

If errors occur when a tape is being read, you can retension the tape, clean the tape drive, and then try again.

### Retensioning a Magnetic Tape Cartridge

Retension a magnetic tape cartridge with the `mt` command.

For example:

```
$ mt -f /dev/rmt/1 retension
$
```

---

**Note** – Do not retension non-QIC tape drives.

---

### Rewinding a Magnetic Tape Cartridge

To rewind a magnetic tape cartridge, use the `mt` command.

For example:

```
$ mt -f /dev/rmt/1 rewind
$
```

## Guidelines for Drive Maintenance and Media Handling

A backup tape that cannot be read is useless. So, periodically clean and check your tape drives to ensure correct operation. See your hardware manuals for instructions on procedures for cleaning a tape drive. You can check your tape hardware by copying some files to the tape, reading the files back, and then comparing the original files with the copied files

Be aware that hardware can fail in ways that the system does not report.

Always label your tapes after a backup. This label should never change. Every time you do a backup, make another tape label that contains the following information:

- The backup date
- The name of the machine and file system that is backed up
- The backup level
- The tape number (1 of  $n$ , if the backup spans multiple volumes)
- Any information specific to your site

Store your tapes in a dust-free safe location, away from magnetic equipment. Some sites store archived tapes in fireproof cabinets at remote locations.

You should create and maintain a log that tracks which media (tape volume) stores each job (backup) and the location of each backed-up file.

# Index

---

## Numbers and Symbols

/export/home directory, 286  
1394 (FireWire) support, description, 111  
4.3 Tahoe file system, 281  
9660 DVD format, 27

## A

accessing  
  disk devices, 61  
  iSCSI disks (how to), 245  
  removable media (how to), 29  
  tape devices, 64  
adding  
  a SCSI device to a SCSI bus (how to), 77  
  a USB camera (how to), 113  
  a USB mass storage device (how to), 112  
  adding a new removable media drive (how to), 28  
  entry to /etc/vfstab file (how to), 302  
  PCI adapter card (how to), 86  
  swap to vfstab, 313  
  USB audio devices (how to), 127  
  VPPA communication service (how to), 149  
autoconfiguration process, 50  
autofs, 289–290  
automounting, and /home, 290

## B

backup, device names, 334–336

block disk device interface  
  definition, 62  
  when to use, 62  
BSD Fat Fast File system, 281  
bus-oriented disk controllers, 63

## C

cdwr command  
  description, 35  
  restricting access to (how to), 38  
  writing data CDs and DVDs and audio CDs  
    (overview), 37  
cfgadm  
  PCI hot-plugging (overview), 66  
  SCSI hot-plugging (overview), 66  
changing, primary USB audio device (how to), 129  
checking, a file system on removable media (how to), 24  
clients  
  iSNS, 258  
  iSNS, display status, CLI, 266  
  iSNS, managing, CLI, 265–268  
configuring  
  a SCSI controller (how to), 74  
  a SCSI device (how to), 74  
  a USB device (how to), 132  
  IB Port, HCA\_SVC, or a VPPA device (how to), 145  
  IB pseudo device (how to), 145  
  iSNS server, 259, 262–265  
  SATA device (how to), 89

- connecting
  - a SCSI controller (how to), 76
  - a USB device, logically (how to), 133
- copying
  - directories between file systems with `cpio` command (overview), 321
  - groups of files with `cpio` command (overview), 321
  - individual files with `cpio` command (overview), 321
- core files, finding and deleting, 187
- `cpio` command
  - (overview), 321
  - copying directories between file systems (how to), 322
  - extract all files from tape (how to), 329
  - listing files on tape (how to), 328
- creating
  - a data CD or DVD file system (how to), 41
  - a disk slice for a ZFS root file system SPARC, 210
  - a file system on a DVD-RAM (how to), 24
  - a file system on a USB mass storage device (how to), 116
  - a file system on removable media (how to), 23
  - a Solaris partition and modifying the slices on a USB mass storage device (how to), 121
  - disk slice for a ZFS root file system (how to) x86, 219
  - loopback file system (overview), 294
- CTFS file system, 284
  
- D**
- `datadm` command, 152
- default
  - file system for `/tmp` (TMPFS), 283
  - SunOS file system, 285
- default discovery domain, iSNS, 258
- default discovery domain set, enabling, 263
- deleting
  - core files, 187
  - finding and deleting old/inactive files, 184
  - temporary files, 187
- detecting end of media, `cpio` command, 321
- determining
  - file system types, 291
  - mounted file systems, 295
- `/dev/dsk` directory, 62
- `/dev/rdisk` directory, 62
- `devfsadm` command, 60
- device driver
  - adding, 59
  - definition, 49
- device names, backup, 334–336
- devices, accessing, 60
- `df` command
  - t option (total blocks), 178
  - examples, 177
  - overview, 176
- directories
  - copying between file systems with `cpio` command (overview), 321
  - displaying information about, 178, 179, 182, 183
  - `/proc`, 283
  - size of, 182, 183
  - temporary, clearing out, 184, 187
  - `/tmp`, 283
- disabling
  - removable media services (how to), 28
  - write protection on removable media (how to), 26
- disconnect, a USB device, logically (how to), 132
- disconnecting
  - a SCSI controller (how to), 75
  - a USB device subtree, logically (how to), 133
- discovery domain
  - client
    - removing, CLI, 267
  - iSNS
    - adding, CLI, 264
    - assigning clients, CLI, 264–265
    - creating, CLI, 263–264
    - displaying status, CLI, 266
    - removing, 267
- discovery domain set
  - disabling
    - displaying status, CLI, 268
  - iSNS
    - creating, CLI, 263

- discovery domain set, iSNS (*Continued*)
    - displaying status, CLI, 266
    - removing, 268
  - disk
    - formatting a (overview), 170
    - setting up for ZFS file systems (overview)
      - x86, 216–229
    - when to format (overview), 192
  - disk-based file systems, 281
  - disk controllers, 63
  - disk for a ZFS file system
    - setting up (how to)
      - SPARC, 214
  - disk label
    - creating (overview), 197
    - description, 163
  - disk setup for a ZFS root file system
    - connecting (how to)
      - SPARC, 209
      - x86, 217
  - disk slice
    - creating for a ZFS root file system
      - SPARC, 210
    - for a ZFS root file system (how to)
      - x86, 219
  - disk slices
    - definition, 167
    - displaying information about (overview), 194
  - disk space
    - displaying information about
      - directory sizes, 182, 183
      - disk space owned per user, 183
      - file sizes, 178, 179, 182
    - finding and deleting old/inactive files, 184, 188
    - finding files exceeding a size limit, 181
    - finding large files, 180
  - diskettes
    - creating a file system on (how to), 23
    - formatting with `rmformat` (how to), 22
    - loading with volume management (how to), 21
  - disks
    - determining if formatted (how to), 193
    - recovering a corrupted disk label (how to), 204
    - recovering a corrupted disk label (overview), 203
  - disks (*Continued*)
    - setting up for ZFS file systems (overview)
      - SPARC, 208
  - displaying
    - directory information, 178, 179, 182
    - disk slice information (overview), 194
    - file information
      - file size, 178, 180
      - listing newest, 184
      - using the `du` command, 182
    - file system information, 183
    - IB communication services (how to), 149
    - InfiniBand device information (how to), 142
    - information about SCSI devices, 73
    - kernel IB clients of an HCA (how to), 146
    - PCI slot configuration information (how to), 83
    - removable media user (how to), 29
    - SATA devices (how to), 88
    - size of files, 179–180
    - swap space, 315–316
    - system configuration information, 51, 55
    - USB bus information (how to), 131
    - USB device information (how to), 115
  - displaying settings, iSNS server, 262
  - DOS, file system, 281
  - driver not attached message, 52
  - `du` command, 182, 183
  - `dumpadm` command, 62
  - DVD-RAM, creating a file system on (how to), 24
  - DVD-ROM, 282
  - DVDs, ISO 9660 format, 27
  - dynamic reconfiguration, InfiniBand devices, 141
  - dynamic reconfiguration (overview), 66
- ## E
- EFI label
    - (overview), 163
    - comparison with VTOC label, 164
    - installing a system with, 166
    - restrictions of, 165
    - troubleshooting problems, 167
  - `eject` command, removable media (how to), 30
  - ejecting, removable media (how to), 30

## enabling

- removable media services (how to), 28
  - uDAPL, 151
  - write protection on removable media (how to), 26
- end-of-media detection, cpio command, 321

**F**

FDFS file system, 284

FIFOFS file system, 284

## file systems

- 4.3 Tahoe, 281
- BSD Fat Fast, 281
- creating (overview)
  - loopback (LOFS), 294
- CTFS, 284
- default SunOS, 285
- disk-based, 281
- DOS, 281
- /export/home, 286
- FDFS, 284
- FIFOFS, 284
- finding types, 291
- High Sierra, 281
- ISO 9660, 281
- making available (overview), 294
- MNTFS, 286
- mount table, 287
- NAMEFS, 284
- network-based, 282
- OBJFS, 284
- /opt, 286
- PCFS, 281
- /proc, 286
- process, (overview), 283–284
- PROCFS, (overview), 283–284
- pseudo, (overview), 282
- SHAREFS, 284
- sharing, 288–289
- SPECFS, 284
- stopping all processes accessing (how to), 306
- SWAPFS, 284
- TMPFS, 283
- types of, 280

file systems (*Continued*)

- UFS, 281
  - UNIX, 281
  - /var, 286
  - ZFS, 281
- files
- deleting
    - See* deleting
  - displaying information about
    - listing, 178, 179
    - size, 178, 179, 182, 183
  - displaying size of, 179–180
  - /etc/default/fs, 291
  - /etc/dfs/fstypes, 291
  - finding files exceeding a size limit, 181
  - in the /proc directory, 284
  - retrieving from tape with tar command (how to), 325
  - sharing, 288–289
  - size of, 178, 179, 182, 183
- find command
- core files, 187
  - finding files exceeding a size limit, 181
  - old/inactive files, 184, 185
- finding
- and deleting old/inactive files
    - See* deleting
  - files exceeding a size limit, 181
  - large files, 180
  - type of file system, 291
- format utility
- (overview), 168
  - analyze menu, 273
  - creating a Solaris fdisk partition (how to), 230
  - defect menu, 275–276
  - determining if a disk is formatted (how to), 193
  - displaying disk slice information (example of), 195
  - entering command names (how to), 276
  - fdisk menu, 272
  - features and benefits, 168
  - guidelines for using, 169–170
  - identifying disks on a system (examples of), 191
  - identifying disks on a system (how to), 190
  - input to, 276, 277

format utility (*Continued*)

- labeling a disk
  - example of, 198
- main menu, 270
- partition menu, 272
- recovering corrupted disk label (how to), 204
- using help facility, 277
- when to use, 169

formatting, diskettes with `rmformat` (how to), 22

formatting a disk, (overview), 170

free hog slice, *See* donor slice

`fstypes` file, 291

## fuser command

- finding if removable media is in use (how to), 29
- killing processes accessing removable media (how to), 29

**H**

High Sierra file system, 281

`/home` (automounted), 290

## hot-plugging

- (overview), 66
- adding a SCSI device to a SCSI bus (how to), 77
- adding PCI adapter card (how to), 86
- configuring a SCSI controller (how to), 74
- configuring a SCSI device (how to), 74
- configuring a USB device (how to), 132
- connecting a SCSI controller (how to), 76
- disconnecting a SCSI controller with `cfgadm` command (how to), 75
- logically connecting a USB device (how to), 133
- logically disconnecting a USB device (how to), 132
- logically disconnecting a USB device subtree (how to), 133
- PCI devices (overview), 82
- removing a SCSI device (how to), 79
- removing PCI adapter card (how to), 84
- replacing an identical SCSI device on a SCSI controller (how to), 78
- unconfiguring a SCSI device (how to), 73
- unconfiguring a USB device (how to), 131

HSFS, *See* High Sierra file system

**I**

## identifying

- devices, 53
- disks on a system (how to), 190
- primary USB audio device (how to), 128

## InfiniBand devices

- adding a VPPA communication service (how to), 149
- configuring an IB Port, HCA\_SVC, or a VPPA device (how to), 145
- configuring an IB pseudo device (how to), 145
- displaying (how to), 142
- displaying IB communication services (how to), 149
- displaying kernel IB clients of an HCA (how to), 146
- dynamic reconfiguration (overview), 141
- overview, 137
- removing an existing IB port, HCA\_SVC, or a VPPA communication service (how to), 150
- unconfiguring an IB Port, HCA\_SVC, or a VPPA (how to), 144
- unconfiguring an IB pseudo device (how to), 145
- updating the IB `P_key` tables (how to), 148

## initiators

- iSNS, display status, CLI, 266
- iSNS, managing, CLI, 265–268

`installboot` command for a ZFS root file system, 214

`installgrub` command for a ZFS root file system, 227

installing boot blocks for a ZFS root file system (how to), SPARC, 214

Internet Storage Name Service, *See* iSNS

## iSCSI

- accessing iSCSI disks (how to), 245
- static and dynamic target discovery, 237

## iSNS

- clients, 258
- default discovery domain, 258
- display client status, 266
- display initiator status, 266
- display target status, 266
- managing clients, 265–268
- managing initiators, 265–268
- managing targets, 265–268
- nodes, 258

**iSNS (Continued)**

- technology overview, 257

**iSNS server**

- configuring, 259

- configuring, CLI, 262–265

**discovery domain**

- adding, CLI, 264

- assigning clients, CLI, 264–265

- creating, CLI, 263–264

- displaying status, CLI, 266

- removing, CLI, 267

- removing a client, 267

**discovery domain set**

- creating, CLI, 263

- disabling, CLI, 268

- displaying status, CLI, 266

- removing, 268

**display settings, 262**

- enable, disable server state notifications, 261

- enabling default discovery domain set, 263

- set retry threshold, 261

- specify data store location, 257, 261

- ISO 9660 file system, 281

- ISO standards, 9660 DVD format, 27

**K**

- /kernel/drv directory, 51

**killing**

- all processes accessing a file system (how to), 306

- processes accessing removable media (how to), 29

**L**

- large files, 180

**listing**

- files and directories, 178, 179, 184, 185

- loading, diskettes (how to), 21

**logical device name**

- definition, 61

- disk, 61

- tape, 64

- logical device names, removable media, 64

- loopback file system (LOFS), creating (overview), 294

**ls command**

- checking directory sizes, 179

- l option (size in bytes), 179

- s option (size in blocks), 180

- t option (newest files), 184

**M**

- maintaining tape drives, 337

**managing**

- iSNS, 266

- clients, CLI, 265–268

- maximums, finding files exceeding maximum size, 181

- media was found message, 21

- memory storage (virtual), definition, 309

- mkisofs command, create a data CD or DVD file

- system (how to), 41

- MNTFS file system, 286

- mnttab file, 287

- modifying, partitions and creating a PCFS file system on

- a USB mass storage device (how to), 118

- mount point, definition, 287

- mount table, 287

**mounting**

- a file system with /etc/vfstab, 303

- a USB mass storage device (how to), 123

- file systems automatically, 289–290

- NFS file systems, 303

- remote removable media manually (example of), 32

- removable media

- automatic mounting compared to, 18

- UFS file systems, 303

- mt command, 337

**N**

- NAMEFS file system, 284

- network-based file systems, 282

**NFS**

- description, 289

- server description, 289

- vfstab entry for, 303

nfsd daemon  
     starting, 31  
     verifying if running, 31  
 no media was found message, 22  
 nodes  
     iSNS, 258  
     iSNS, display status, CLI, 266  
     iSNS, managing, CLI, 265–268

## O

OBJFS file system, 284  
 /opt directory, 286  
 overview, iSNS technology, 257

## P

partition (swap), definition, 309  
 PCFS file system, 281  
 PCI devices  
     adding PCI adapter card (how to), 86  
     displaying PCI slot configuration information (how to), 83  
     removing PCI adapter card (how to), 84  
     troubleshooting PCI configuration problems, 88  
 physical device name  
     definition, 61  
 /proc directory, 283, 286  
 process file system (PROCFS), 283–284  
 PROCFS file system, (overview), 283–284  
 prtvtoc command, 62  
     (example of), 202  
 pseudo file systems, (overview), 282

## Q

quot command, 183

## R

raw disk device interface, 62

RCM script  
     commands for, 91  
     overview, 90  
 registering, service provider in the DAT static registry (how to), 152  
 removable media  
     accessing (how to), 29  
     accessing media on other systems (example of), 32  
     adding a new removable media drive (how to), 28  
     checking a file system on (how to), 24  
     creating a file system on (how to), 23  
     creating a file system on a DVD-RAM (how to), 24  
     disabling or enabling removable media services (how to), 28  
     ejecting (how to), 30  
     enabling write protection on (how to), 26  
     finding out if media is in use (how to), 29  
     formatting diskettes with rmformat (how to), 22  
     killing processes accessing (how to), 29  
     loading diskettes (how to), 21  
     management, benefits, 17  
     /media  
         mount point, 27  
     mounting  
         manual compared to automatic, 18  
         mounting remote media (example of), 32  
         repairing bad blocks on removable media (how to), 25  
 removing  
     a SCSI device (how to), 79  
     a USB mass storage device (how to), 114  
     existing IB port, HCA\_SVC, or a VPPA communication service (how to), 150  
     PCI adapter card (how to), 84  
 repairing, bad blocks on removable media (how to), 25  
 replacing  
     a disk for a ZFS root pool (how to)  
         x86, 222, 224  
     an identical SCSI device on a SCSI controller (how to), 78  
 resetting, a USB device (how to), 134  
 resolving, a failed SCSI unconfigure operation (how to), 81  
 restricting, removable media access (how to), 38

retrieving, files from tape with tar command (how to), 325

rm command, 185, 187

Rock Ridge extension (HSFS file system), 281

root (/) file system or directory, 286

## S

SATA devices

configuring SATA device (how to), 89

displaying SATA device information (how to), 88

unconfiguring SATA device (how to), 88

SCSI devices

adding a SCSI device to a SCSI bus (how to), 77

configuring a SCSI controller (how to), 74

configuring a SCSI device (how to), 74

connecting a SCSI controller (how to), 76

disconnecting with cfdisk command (how to), 75

displaying information about (how to), 73

removing a SCSI device (how to), 79

replacing an identical SCSI device on a SCSI controller (how to), 78

resolving a failed SCSI unconfigure operation (how to), 81

troubleshooting SCSI configuration problem, 80

unconfiguring a SCSI controller (how to), 73

SCSI tape drives, 335

setting up

a disk for ZFS file systems (overview)  
x86, 216–229

disks for ZFS file systems (overview)  
SPARC, 208

share command, 289

making removable media available to other systems  
(how to), 31

shareall command, 289

SHAREFS file system, 284

sharing, files, 288–289

size

directory, 182, 183

file, 178, 179, 182, 183

slice (definition), 167

Solaris fdisk partition, guidelines, 229–230

SPECFS file system, 284

specifying a disk slice, 63

starting, nfsd daemon, 31

stopping

all processes for a file system (how to), 306

killing processes accessing removable media (how to), 29

storage (virtual memory), definition, 309

storage capacities (media), 333

SunOS default file system, 285

swap command, 63

swap file

adding to vfstab, 313

displaying, 315–316

swap partition, definition, 309

swapadd command, 313

SWAPFS file system, 284

sysdef command, 53

## T

tape

retrieving files from with tar command (how to), 325

sizes, 333

storage capacities, 333

tape devices (naming), 64

tape drive

maintaining, 337

maximum SCSI, 335

rewind, 335

tar command

(overview), 323

copying files to remote tape with dd command (how to), 330

listing files on tape (how to), 324

retrieving files from remote tape with dd command  
(how to), 331

retrieving files from tape (how to), 325

targets

iSNS, display status, CLI, 266

iSNS, managing, CLI, 265–268

temporary directories, 184, 187

temporary file system (TMPFS), overview, 283

/tmp directory, 283, 286

TMPFS file system, overview, 283  
 troubleshooting  
   a failed SCSI unconfigure operation, 81  
   EFI disk labels, 167  
   PCI configuration problems, 88  
   SCSI configuration problems, 80  
   USB audio device problems, 129  
   USB mass storage devices, 124  
 type of file systems, 280

## U

uDAPL  
 (overview), 150  
 enabling (how to), 151  
 registering a service provider in the DAT static registry (how to), 152  
 unregistering a service provider in the DAT static registry (how to), 152  
 updating the DAT static registry (how to), 152  
 UDF file system, 281  
 UFS file system, 281  
   mounting, 303  
   mounting with `/etc/vfstab`, 303  
 UFS file systems, displaying information about, 183  
 unconfiguring  
   a SCSI controller (how to), 73  
   a USB device (how to), 131  
   IB Port, HCA\_SVC, or a VPPA Device (how to), 144  
   IB pseudo device (how to), 145  
   SATA device (how to), 88  
 UNIX file system, 281  
 unmounting  
   a USB mass storage device (how to), 123  
 unregistering, service provider in the DAT static registry (how to), 152  
 unsupported devices, 51  
 updating  
   DAT static registry (how to), 152  
   IB P\_key tables (how to), 148  
 USB devices  
   (overview), 105  
   acronyms, 105  
   USB devices (*Continued*)  
     adding a USB camera (how to), 113  
     adding a USB mass storage device (how to), 112  
     adding audio devices (how to), 127  
     audio  
       (overview), 126  
       changing the primary device (how to), 129  
       device ownership, 129  
     bus description, 107  
     bus-powered devices, 102  
     cables for, 105  
     composite device, 108  
     compound device, 108  
     configuring a USB device (how to), 132  
     connect a USB device (how to), 133  
     creating a file system on a mass storage device (how to), 116  
     creating a Solaris partition and modifying the slices on a mass storage device (how to), 121  
     device classes, 108  
     device nodes, 106  
     diskette devices (overview), 110  
     displaying bus information (how to), 131  
     displaying USB device information (how to), 115  
     drivers, 108  
     host controller and root hub, 103  
     hot-plugging (overview), 112  
     hotpluggable attribute, 111  
     identifying primary audio device (how to), 128  
     keyboards and mouse devices, 102  
     logically disconnecting a USB device (how to), 132  
     logically disconnecting a USB device subtree (how to), 133  
     modifying partitions and creating a PCFS file system on a mass storage device (how to), 118  
     mounting mass storage (how to), 123  
     mounting or unmounting a mass storage device (how to), 123  
     names of, 108  
     Oracle Solaris USB Architecture (USBA), 106  
     physical device hierarchy, 107  
     power management, 104  
     removable mass storage (overview), 109  
     removing a USB mass storage device (how to), 114

USB devices (*Continued*)

- resetting a USB device (how to), 134
  - troubleshooting audio device problems, 129
  - troubleshooting tips for mass storage devices, 124
  - unconfiguring a device (how to), 131
  - unmounting mass storage (how to), 123
  - USB 2.0 support, 102
  - wheel mouse support, 103
- user ownership of disk space, 183

**V**

- /var directory, 286
- /var/adm directory, controlling size of, 186
- verifying, nfsd daemon is running, 31
- vfstab file, 291, 313
  - adding entries to (how to), 302
  - adding swap to, 313
  - entry for LOFS, 302
- virtual memory storage, definition, 309

**W**

- writing, data CDs and DVDs and audio CDs
  - (overview), 37

**Z**

- ZFS file system, 281
  - setting up a disk (how to)
    - x86, 228
- ZFS root pool
  - replacing a disk (how to)
    - x86, 222, 224