

System Administration Guide: Basic Administration

Copyright © 1998, 2011, 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 RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 RIGHTS. Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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 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.

AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. 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. UNIX est une marque déposée concédée sous licence par X/Open Company, Ltd.

Contents

Preface	19
1 Oracle Solaris Management Tools (Road Map)	23
What's New in Oracle Solaris Management Tools?	23
Matrix of Oracle Solaris Management Tools and Supported Releases	25
Feature Descriptions for Oracle Solaris Management Tools	25
Feature Descriptions for Solaris 9 Management Tools	26
Availability of Solaris Management Commands	27
Solaris 10 System Management Commands	27
For More Information About Oracle Solaris Management Tools	29
2 Working With the Solaris Management Console (Tasks)	31
Solaris Management Console (Overview)	31
What Is the Solaris Management Console?	31
Solaris Management Console Tools	32
Why Use the Solaris Management Console?	34
Organization of the Solaris Management Console	34
Changing the Solaris Management Console Window	36
Solaris Management Console Documentation	36
How Much Role-Based Access Control?	36
Becoming Superuser (root) or Assuming a Role	37
▼ How to Become Superuser (root) or Assume a Role	38
Using the Solaris Management Tools With RBAC (Task Map)	40
If You Are the First to Log In to the Console	41
Creating the Primary Administrator Role	41
▼ How to Create the First Role (Primary Administrator)	42
▼ How to Assume the Primary Administrator Role	43

Starting the Solaris Management Console	44
▼ How to Start the Console as Superuser or as a Role	44
Using the Oracle Solaris Management Tools in a Name Service Environment (Task Map)	45
RBAC Security Files	46
Prerequisites for Using the Solaris Management Console in a Name Service Environment	47
Management Scope	47
/etc/nsswitch.conf File	48
▼ How to Create a Toolbox for a Specific Environment	48
▼ How to Add a Tool to a Toolbox	50
▼ How to Start the Solaris Management Console in a Name Service Environment	51
Adding Tools to the Solaris Management Console	51
▼ How to Add a Legacy Tool to a Toolbox	51
▼ How to Install an Unbundled Tool	52
Troubleshooting the Solaris Management Console	52
▼ How to Troubleshoot the Solaris Management Console	53
3 Working With the Oracle Java Web Console (Tasks)	55
What's New in Administering the Oracle Java Web Console?	55
Oracle Java Web Console Server Management	55
Applications That Are Available to the Oracle Java Web Console	56
Oracle Java Web Console (Overview)	56
What Is the Oracle Java Web Console?	57
Oracle Java Web Console Management Commands	58
Supported Web Browsers	58
Getting Started With the Oracle Java Web Console (Task Map)	58
Getting Started With the Oracle Java Web Console	59
▼ How to Start Applications From the Oracle Java Web Console's Launch Page	60
Managing the Console Service	62
▼ How to Start the Console Service	62
▼ How to Enable the Console Service to Run at System Start	62
▼ How to Stop the Console Service	63
▼ How to Disable the Console Service	63
Configuring the Oracle Java Web Console	64
▼ How to Change Oracle Java Web Console Properties	66

Oracle Java Web Console User Identity	68
Using the Console Debug Trace Log	69
Troubleshooting the Oracle Java Web Console Software (Task Map)	70
Troubleshooting the Oracle Java Web Console Software	71
Checking Console Status and Properties	71
Problems Accessing the Console	73
Problems with Application Registration	74
Oracle Java Web Console Reference Information	78
Oracle Java Web Console Security Considerations	78
Specifying Authorizations With the authTypes Tag	80
Enabling Remote Access to the Oracle Java Web Console	82
Disabling Remote Access to the Oracle Java Web Console	82
Changing Internal Passwords for Oracle Java Web Console	83
4 Managing User Accounts and Groups (Overview)	85
What's New or Changed in Managing Users and Groups?	85
Tools for User Account and Group Account Management	86
What Are User Accounts and Groups?	86
User Account Components	87
Guidelines for Using User Names, User IDs, and Group IDs	93
Where User Account and Group Information Is Stored	93
Fields in the passwd File	94
Default passwd File	94
Fields in the shadow File	95
Fields in the group File	96
Default group File	96
Tools for Managing User Accounts and Groups	98
Tasks for Solaris User and Group Management Tools	98
Managing Users and Resources With Projects	101
Customizing a User's Work Environment	101
Customizing the Bash Shell	103
Using Site Initialization Files	103
Avoiding Local System References	103
Shell Features	104
Shell Environment	105

The PATH Variable	107
Locale Variables	108
Default File Permissions (umask)	109
User and Site Initialization Files Examples	110
5 Managing User Accounts and Groups (Tasks)	113
Setting Up User Accounts (Task Map)	113
Gathering User Information	114
▼ How to Customize User Initialization Files	115
▼ How to Add a Group With the Solaris Management Console's Groups Tool	116
▼ How to Add a User With the Solaris Management Console's Users Tool	117
Adding Groups and Users With Command-Line Tools	119
Setting Up Home Directories With the Solaris Management Console	120
▼ How to Share a User's Home Directory	120
▼ How to Mount a User's Home Directory	121
Maintaining User Accounts (Task Map)	123
Modifying User Accounts	123
▼ How to Modify a Group	125
▼ How to Delete a Group	125
Administering Passwords	126
Using Password Aging	127
▼ How to Disable a User Account	127
▼ How to Change a User's Password	128
▼ How to Set Password Aging on a User Account	129
▼ How to Delete a User Account	129
6 Managing Client-Server Support (Overview)	131
What's New in Managing Client-Server Support?	132
Support for Specifying Platform by Using the <code>bootadm -p</code> Command	132
<code>nfs4_domain</code> Keyword Impacts Diskless Client Boot	132
x86: Diskless Client Changes That Apply to GRUB	132
x86: Changes to the <code>smdiskless</code> Command	133
Where to Find Client-Server Tasks	133
What Are Servers, Clients, and Appliances?	134
What Does Client Support Mean?	134

Overview of System Types	135
Description of a Server	135
Stand-Alone Systems	136
Diskless Clients	136
Description of an Appliance	137
Guidelines for Choosing System Types	137
Diskless Client Management Overview	138
OS Server and Diskless Client Support Information	138
Diskless Client Management Features	139
Disk Space Requirements for OS Servers	142
7 Managing Diskless Clients (Tasks)	143
Managing Diskless Clients (Task Map)	143
Preparing for Managing Diskless Clients	145
▼ x86: How to Prepare for Adding Diskless Clients in a GRUB Based Boot Environment ..	146
▼ How to Prepare for Adding Diskless Clients in Oracle Solaris 10	148
▼ How to Add OS Services for Diskless Client Support	150
▼ x86: How to Add a Diskless Client in the GRUB Based Boot Environment	152
▼ How to Add a Diskless Client in Oracle Solaris 10	155
▼ x86: How to Boot a Diskless Client With GRUB	157
▼ SPARC: How to Boot a Diskless Client in Oracle Solaris 10	158
▼ How to Remove Diskless Client Support	158
▼ How to Remove OS Services for Diskless Clients	159
Patching Diskless Client OS Services	160
Displaying OS Patches for Diskless Clients	160
▼ How to Add an OS Patch for a Diskless Client	161
Troubleshooting Diskless Client Problems	163
Troubleshooting Diskless Client Installation Problems	163
Troubleshooting General Diskless Client Problems	167
8 Introduction to Shutting Down and Booting a System	171
What's New in Shutting Down and Booting a System	171
Support for Fast Reboot on the SPARC Platform	172
Oracle Solaris Auto Registration Introduced	172
Automatic Boot Archive Recovery	173

SPARC Support for Install-Time Updates	173
Two-Terabyte Disk Support for Installing and Booting Oracle Solaris 10	174
Oracle Solaris ZFS Boot Support	174
x86: findroot Command	174
Support for Specifying Platform by Using bootadm Command	175
Redesign of SPARC Bootstrap Process	175
x86: Support for Using Power Button to Initiate System Shutdown	176
Where to Find Shut Down and Boot Tasks	176
Shut Down and Boot Terminology	177
Guidelines for Shutting Down a System	177
Guidelines for Booting a System	178
When to Shut Down a System	179
When to Boot a System	180
9 Shutting Down and Booting a System (Overview)	183
Fundamentals of the Oracle Solaris Boot Design	184
Understanding the New SPARC Boot Architecture	185
Packing and Unpacking the Miniroot	186
Software Installation and Upgrades	186
Installation Memory Requirements	187
Changes to the Network Boot Server Setup Process	187
Support for Booting Multiple Kernels	187
Implementation of the Boot Archives on SPARC	187
x86: Administering the GRUB Bootloader	188
How GRUB Based Booting Works	188
GRUB Support for findroot Command	189
Bootting From an Oracle Solaris ZFS Root File System	190
Installation Requirements for Oracle Solaris ZFS	190
How Booting From an Oracle Solaris ZFS Root File System Works	191
SPARC: Options That Support Booting From an Oracle Solaris ZFS Root File System	191
x86: Options That Support Booting From a ZFS Root File System	192
10 Shutting Down a System (Tasks)	193
Shutting Down the System (Task Map)	193
Shutting Down the System	194

System Shutdown Commands	194
User Notification of System Down Time	195
▼ How to Determine Who Is Logged in to a System	196
▼ How to Shut Down a Server	196
▼ How to Shut Down a Stand-Alone System	199
Turning Off Power to All Devices	201
▼ How to Turn Off Power to All Devices	201
11 Modifying Oracle Solaris Boot Behavior (Tasks)	203
Modifying Boot Behavior on SPARC Based Systems (Task Map)	203
SPARC: Using the Boot PROM	204
▼ SPARC: How to Find the PROM Revision Number for a System	205
▼ SPARC: How to Identify Devices on a System	205
▼ SPARC: How to Determine the Default Boot Device	207
▼ SPARC: How to Change the Default Boot Device by Using the Boot PROM	207
▼ SPARC: How to Change the Default Boot Device by Using the eeprom Command	209
SPARC: Resetting the System	209
▼ SPARC: How to Change the Default Kernel by Using the Boot PROM	209
▼ SPARC: How to Change the Default Kernel by Using the eeprom Command	210
SPARC Support for ITU Construction Tools	210
Modifying Boot Behavior on x86 Based Systems (Task Map)	212
Modifying Boot Behavior on x86 Based Systems	213
x86: Modifying Boot Behavior by Editing the GRUB Menu at Boot Time	214
Editing the GRUB Menu in Oracle Solaris 10	215
Boot Arguments That Can Be Specified by Editing the GRUB Menu at Boot Time	216
▼ x86: How to Modify Boot Behavior by Editing the GRUB Menu at Boot Time	217
x86: Modifying Boot Behavior by Editing the menu.lst File	218
▼ x86: How to Modify Boot Behavior by Editing the menu.lst File	219
x86: Locating the Active GRUB menu.lst File	222
x86: Implementation of the findroot Command	223
▼ x86: How to Add GRUB Menu Entries That Use the findroot Command	224
12 Booting an Oracle Solaris System (Tasks)	227
Booting a SPARC Based System (Task Map)	228
Booting a SPARC Based System	229

▼ SPARC: How to Boot a System to Run Level 3 (Multiuser Level)	229
▼ SPARC: How to Boot a System to Run Level S (Single-User Level)	230
▼ SPARC: How to Boot a System Interactively	231
▼ SPARC: How to Boot a Kernel Other Than the Default Kernel	232
Booting From a Specified ZFS Root File System on a SPARC Based System	234
▼ SPARC: How to List Available Bootable Datasets Within a ZFS Root Pool	235
▼ SPARC: How to Boot From a Specified ZFS Root File System	236
Booting a SPARC Based System in Failsafe Mode	238
▼ How to Boot a SPARC Based System in Failsafe Mode	239
Booting a SPARC Based System From the Network	243
▼ SPARC: How to Boot a System From the Network	243
Booting an x86 Based System by Using GRUB (Task Map)	245
▼ x86: How to Boot a System to Run Level 3 (Multiuser)	246
▼ x86: How to Boot a System to Run Level S (Single-User Level)	247
▼ x86: How to Boot a System Interactively	249
x86: Booting From a Specified ZFS Root File System on an x86 Based System	252
▼ x86: How to Display a List of the Available ZFS Boot Environments	252
▼ x86: How to Boot From a Specified ZFS Root File System	253
Booting an x86 Based System in Failsafe Mode	255
▼ How to Boot an x86 Based System in Failsafe Mode	256
▼ x86: How to Boot in Failsafe Mode to Forcibly Update a Corrupt Boot Archive	258
Booting an x86 Based System From the Network	260
x86: About DHCP Macros	261
▼ x86: How to Perform a GRUB Based Boot From the Network	263
Accelerating the Reboot Process on the SPARC Platform (Task Map)	264
Initiating a Fast Reboot of a SPARC Based System	265
▼ How to Initiate a Fast Reboot of a SPARC Based System	265
Performing a Standard Reboot of a SPARC Based System	265
Managing the Boot Configuration Service	265
13 Managing the Oracle Solaris Boot Archives (Tasks)	267
Managing the Oracle Solaris Boot Archives (Task Map)	267
Description of the Oracle Solaris Boot Archives	268
Managing the boot - archive Service	269
▼ How to Enable or Disable the boot - archive Service	269

Automatic Boot Archive Recovery	270
▼ x86: How to Clear Automatic Boot Archive Update Failures by Using the auto-reboot-safe Property	270
▼ How to Clear Automatic Boot Archive Update Failures by Using the bootadm Command	271
Using the bootadm Command to Manage the Boot Archives	272
▼ How to Manually Update the Boot Archive by Using the bootadm Command	272
▼ How to Manually Update the Boot Archive on a Solaris Volume Manager RAID-1 (Mirrored) Root Partition	273
▼ How to List Contents of the Boot Archive	275
▼ x86: How to Locate the Active GRUB Menu and List Current Menu Entries	275
▼ x86: How to Set the Default Boot Entry for the Active GRUB Menu	276
14 Troubleshooting Booting an Oracle Solaris System (Tasks)	277
Troubleshooting Booting on the SPARC Platform (Task Map)	277
▼ SPARC: How to Stop the System for Recovery Purposes	278
SPARC: Forcing a Crash Dump and Reboot of the System	278
▼ SPARC: How to Boot a System for Recovery Purposes	280
▼ SPARC: How to Boot to a ZFS Root Environment to Recover From a Lost Password or Similar Problem	282
▼ SPARC: How to Boot the System With the Kernel Debugger (kldb)	283
Troubleshooting Booting on the x86 Platform (Task Map)	284
▼ x86: How to Stop a System for Recovery Purposes	284
x86: Forcing a Crash Dump and Reboot of the System	284
▼ x86: How to Boot a System With the Kernel Debugger in the GRUB Boot Environment (kldb)	286
15 x86: GRUB Based Booting (Reference)	289
x86: Boot Processes	289
x86: System BIOS	289
x86: Kernel Initialization Process	290
x86: Support for GRUB in the Oracle Solaris OS	290
x86: GRUB Terminology	290
x86: Functional Components of GRUB	292
How Multiple Operating Systems Are Supported by GRUB	293
x86: Supported GRUB Versions	295

16	x86: Booting a System That Does Not Implement GRUB (Tasks)	297
	x86: Booting a System (Task Map)	297
	x86: Booting a System That Does Not Implement GRUB	298
	▼ x86: How to Boot a System to Run Level 3 (Multiuser Level)	299
	▼ x86: How to Boot a System to Run Level S (Single-User Level)	301
	▼ x86: How to Boot a System Interactively	303
	x86: Booting From the Network	305
	▼ x86: How to Boot a System From the Network	305
	x86: Using the Device Configuration Assistant	307
	▼ x86: How to Stop a System for Recovery Purposes	308
	▼ x86: How to Boot a System for Recovery Purposes	308
	▼ x86: How to Boot a System With the Kernel Debugger (kmdb)	311
	x86: Forcing a Crash Dump and Reboot of the System	313
	x64: Troubleshooting a Failed 64-Bit Boot	315
	x86: Boot Processes (Reference)	315
	x86: Boot Subsystems	315
	x86: Boot Process	321
	x86: Boot Files	322
17	Working With the Oracle Solaris Auto Registration <code>regadm</code> Command (Tasks)	325
	Oracle Solaris Auto Registration (Overview)	326
	How Your Oracle Solaris 10 System Is Registered	326
	When and How Service Tag Data Is Transmitted	327
	Managing Auto Registration Configuration	328
	How the Auto Registration Process Works	328
	Getting Started With the <code>regadm</code> Command	330
	When to Use the <code>regadm</code> Command-Line Interface	330
	Requirements for Using the <code>regadm</code> Command	331
	Managing the Auto Registration SMF Service	331
	▼ How to Determine If the Auto Registration SMF Service Is Enabled	332
	▼ How to Disable the Auto Registration SMF Service	332
	▼ How to Enable the Auto Registration SMF Service	332
	Configuring the Auto Registration Feature by Using the <code>regadm</code> Command	333
	Tasks That Are Required Prior to Registration	334
	Displaying the Current Auto Registration Configuration	334

▼ How to Clear the Current Auto Registration Configuration	335
Configuring Auto Registration Properties	336
Authenticating a User by Specifying Named Support Credentials	338
Registering Your System by Using the <code>regadm</code> Command	341
▼ How to Display Registration Information About Products Not Yet Registered	341
▼ How to Register Your System by Using the <code>regadm</code> Command	342
Obtaining Information About Currently Registered Products	344
Troubleshooting Auto Registration Issues	345
Authentication Failed	345
Cannot Send HTTP Proxy Password	345
HTTP Communications Error	345
18 Managing Services (Overview)	347
Introduction to SMF	347
Changes in Behavior When Using SMF	348
SMF Concepts	349
SMF Service	349
Service Identifiers	350
Service States	351
SMF Manifests	351
SMF Profiles	352
Service Configuration Repository	352
SMF Repository Backups	353
SMF Snapshots	353
SMF Administrative and Programming Interfaces	353
SMF Command-Line Administrative Utilities	354
Service Management Configuration Library Interfaces	354
SMF Components	354
SMF Master Restarter Daemon	354
SMF Delegated Restarters	355
SMF and Booting	355
SMF Compatibility	356
Run Levels	356
When to Use Run Levels or Milestones	357
Determining a System's Run Level	357

/etc/inittab File	358
What Happens When the System Is Brought to Run Level 3	359
19 Managing Services (Tasks)	361
Managing Services (Task Map)	361
Monitoring SMF Services	362
▼ How to List the Status of a Service	362
▼ How to Show Which Services Are Dependent on a Service Instance	364
▼ How to Show Which Services a Service Is Dependent On	364
Managing SMF Services (Task Map)	365
Managing SMF Services	365
Using RBAC Rights Profiles With SMF	365
▼ How to Disable a Service Instance	366
▼ How to Enable a Service Instance	366
▼ How to Restart a Service	367
▼ How to Restore a Service That Is in the Maintenance State	367
▼ How to Revert to Another SMF Snapshot	368
▼ How to Create an SMF Profile	369
▼ How to Apply an SMF Profile	370
▼ Changing Services Offered to the Network with generic*.xml	371
Configuring SMF Services	371
▼ How to Modify a Service	371
▼ How to Change an Environment Variable for a Service	372
▼ How to Change a Property for an inetd Controlled Service	373
▼ How to Modify a Command-Line Argument for an inetd Controlled Service	374
▼ How to Convert inetd.conf Entries	375
Using Run Control Scripts (Task Map)	376
Using Run Control Scripts	376
▼ How to Use a Run Control Script to Stop or Start a Legacy Service	376
▼ How to Add a Run Control Script	377
▼ How to Disable a Run Control Script	378
Troubleshooting the Service Management Facility	379
▼ Debugging a Service That Is Not Starting	379
▼ How to Repair a Corrupt Repository	379
▼ How to Boot Without Starting Any Services	382

▼ How to Force a <code>su</code> login Prompt If the <code>system/filesystem/local:default</code> Service Fails During Boot	382
20 Managing Software (Overview)	385
What's New in Software Management in the Oracle Solaris Operating System?	386
Oracle Solaris Auto Registration	386
Package and Patching Tools Enhancements to Support Oracle Solaris Zones	386
Deferred-Activation Patching	387
Common Agent Container Included in the Oracle Solaris OS	387
Improvements to How <code>patchadd -M</code> Command Handles Multiple Patches	388
Package and Patch Tool Enhancements	388
Where to Find Software Management Tasks	389
Overview of Software Packages	389
Tools for Managing Software Packages	390
Adding or Removing a Software Package (<code>pkgadd</code>)	391
Key Points for Adding Software Packages (<code>pkgadd</code>)	391
Guidelines for Removing Packages (<code>pkgrm</code>)	392
Restrictions on Adding and Removing Software Packages and Patches for Solaris Releases That Are Not Zones Aware	392
Avoiding User Interaction When Adding Packages (<code>pkgadd</code>)	393
Using an Administration File	393
Using a Response File (<code>pkgadd</code>)	394
21 Managing Software With Oracle Solaris System Administration Tools (Tasks)	395
Oracle Solaris Product Registry and Solaris GUI Installation Tools for Managing Software ..	395
Adding Software With the Oracle Solaris Installation GUI	396
▼ How to Install Software With the Oracle Solaris Installation GUI Program	396
Managing Software With the Oracle Solaris Product Registry GUI (Task Map)	397
▼ How to View Installed or Uninstalled Software Information With the Oracle Solaris Product Registry GUI	398
▼ How to Install Software With the Oracle Solaris Product Registry GUI	399
▼ How to Uninstall Software With the Oracle Solaris Product Registry GUI	400
Managing Software With the Oracle Solaris Product Registry Command-Line Interface (Task Map)	401
Managing Software With the Oracle Solaris Product Registry Command-Line Interface	402
▼ How to View Installed or Uninstalled Software Information (<code>prodreg</code>)	402

▼ How to View Software Attributes (prodreg)	405
▼ How to Check for Software Dependencies (prodreg)	407
▼ How to Identify Damaged Software Products (prodreg)	408
▼ How to Uninstall Software (prodreg)	410
▼ How to Uninstall Damaged Software (prodreg)	412
▼ How to Reinstall Damaged Software Components (prodreg)	414
22 Managing Software by Using Oracle Solaris Package Commands (Tasks)	417
Managing Software Packages by Using Package Commands (Task Map)	417
Using Package Commands to Manage Software Packages	418
▼ How to Add Software Packages (pkgadd)	418
Adding a Software Package to a Spool Directory	421
▼ How to List Information About All Installed Packages (pkginfo)	423
▼ How to Check the Integrity of Installed Software Packages (pkgchk)	424
▼ How to Check the Integrity of Installed Objects (pkgchk -p, pkgchk -P)	425
Removing Software Packages	427
▼ How to Remove Software Packages (pkgrm)	427
23 Managing Patches	429
About Patches	429
Patching Strategy	430
Live Upgrade	430
Applying an Oracle Solaris Update or an Oracle Solaris Update Patch Bundle	431
Applying a Recommended Patch Cluster	431
Applying a Critical Patch Update	432
Applying an Enterprise Installation Standards Patch Baseline	432
Downloading a Patch	432
▼ How to Search for a Patch	432
Displaying Information About Patches	433
Applying a Patch	433
▼ How to Apply a Patch Using the patchadd Command	434
Removing a Patch	434
Patch Management Terms and Definitions	435

A SMF Services437

Index 443

Preface

System Administration Guide: Basic Administration is part of a documentation 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 10 operating system (OS)
- Set up all the networking software that you plan to use

For the Oracle Solaris 10 release, new features that might be interesting 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 OS: Hardware Compatibility Lists*. This document cites any implementation differences between the platform types.

In this document, these x86 related terms mean the following:

- x86 refers to the larger family of 64-bit and 32-bit x86 compatible products.
- x64 relates specifically to 64-bit x86 compatible CPUs.
- "32-bit x86" points out specific 32-bit information about x86 based systems.

For supported systems, see the *Oracle Solaris OS: Hardware Compatibility Lists*.

Who Should Use This Book

This book is intended for anyone responsible for administering one or more systems running the Oracle Solaris 10 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.

How the System Administration Guides Are Organized

Here is a list of the topics that are covered by the System Administration Guides.

Book Title	Topics
<i>System Administration Guide: Basic Administration</i>	User accounts and groups, server and client support, shutting down and booting a system, managing services, and managing software (packages and patches)
<i>System Administration Guide: Advanced Administration</i>	Terminals and modems, system resources (disk quotas, accounting, and crontabs), system processes, and troubleshooting Oracle Solaris software problems
<i>System Administration Guide: Devices and File Systems</i>	Removable media, disks and devices, file systems, and backing up and restoring data
<i>System Administration Guide: IP Services</i>	TCP/IP network administration, IPv4 and IPv6 address administration, DHCP, IPsec, IKE, Solaris IP filter, Mobile IP, IP network multipathing (IPMP), and IPQoS
<i>System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)</i>	DNS, NIS, and LDAP naming and directory services, including transitioning from NIS to LDAP and transitioning from NIS+ to LDAP
<i>System Administration Guide: Naming and Directory Services (NIS+)</i>	NIS+ naming and directory services
<i>System Administration Guide: Network Services</i>	Web cache servers, time-related services, network file systems (NFS and autofs), mail, SLP, and PPP
<i>System Administration Guide: Printing</i>	Printing topics and tasks, using services, tools, protocols, and technologies to set up and administer printing services and printers
<i>System Administration Guide: Security Services</i>	Auditing, device management, file security, BART, Kerberos services, PAM, Solaris Cryptographic Framework, privileges, RBAC, SASL, and Solaris Secure Shell
<i>System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones</i>	Resource management topics projects and tasks, extended accounting, resource controls, fair share scheduler (FSS), physical memory control using the resource capping daemon (rcapd), and resource pools; virtualization using Solaris Zones software partitioning technology and lx branded zones
<i>Oracle Solaris ZFS Administration Guide</i>	ZFS storage pool and file system creation and management, snapshots, clones, backups, using access control lists (ACLs) to protect ZFS files, using ZFS on an Oracle Solaris system with zones installed, emulated volumes, and troubleshooting and data recovery

Book Title	Topics
<i>Oracle Solaris Trusted Extensions Administrator's Procedures</i>	System administration that is specific to the Oracle Solaris' Trusted Extensions feature
<i>Oracle Solaris Trusted Extensions Configuration Guide</i>	Starting with the Solaris 10 5/08 release, describes how to plan for, enable, and initially configure the Oracle Solaris' Trusted Extensions feature

Related Third-Party Web Site References

Note – Oracle is not responsible for the availability of third-party web sites mentioned in this document. Oracle does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Oracle will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

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.

Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen 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 onscreen computer output	<code>machine_name% su</code> Password:
<i>aabbcc123</i>	Placeholder: replace with a real name or value	The command to remove a file is <code>rm filename</code> .

TABLE P-1 Typographic Conventions (Continued)

Typeface	Meaning	Example
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized	Read Chapter 6 in the <i>User's Guide</i> . <i>A cache</i> is a copy that is stored locally. Do <i>not</i> save the file. Note: Some emphasized items appear bold online.

Shell Prompts in Command Examples

The following table shows the default UNIX system prompt and superuser prompt for shells that are included in the Oracle Solaris OS. Note that the default system prompt that is displayed in command examples varies, depending on the Oracle Solaris release.

TABLE P-2 Shell Prompts

Shell	Prompt
Bash shell, Korn shell, and Bourne shell	\$
Bash shell, Korn shell, and Bourne shell for superuser	#
C shell	machine_name%
C shell for superuser	machine_name#

General Conventions

Be aware of the following conventions used in this book.

- When following steps or using examples, be sure to type double-quotes ("), left single-quotes ('), and right single-quotes (') exactly as shown.
- The key referred to as Return is labeled Enter on some keyboards.
- The root path usually includes the `/sbin`, `/usr/sbin`, `/usr/bin`, and `/etc` directories, so the steps in this book show the commands in these directories without absolute path names. Steps that use commands in other, less common, directories show the absolute paths in the examples.

Oracle Solaris Management Tools (Road Map)

This chapter provides a roadmap to Solaris management tools.

- “What's New in Oracle Solaris Management Tools?” on page 23
- “Matrix of Oracle Solaris Management Tools and Supported Releases” on page 25
- “Feature Descriptions for Oracle Solaris Management Tools” on page 25
- “Feature Descriptions for Solaris 9 Management Tools” on page 26
- “For More Information About Oracle Solaris Management Tools” on page 29

What's New in Oracle Solaris Management Tools?

These tools are new or changed as of the Oracle Solaris 10 initial 3/05 release:

- `admintool` – Starting with Oracle Solaris 10, this tool is no longer available
- Package and Patch Tool Enhancements

For a complete listing of new features and a description of Oracle Solaris releases, see [Oracle Solaris 10 8/11 What's New](#).

The following table provides a brief description of new or changed management tools.

TABLE 1-1 New or Changed Management Tools in the Oracle Solaris Release

Solaris Management Tool	Description	For More Information
admintool	<p>This tool is no longer available.</p> <p>Alternative tools include the following:</p> <ul style="list-style-type: none"> ■ Solaris Management Console to manage users, groups, terminals, and modems ■ Oracle Solaris Product Registry to manage software ■ Solaris Print Manager to manage printers 	<p>“Setting Up User Accounts (Task Map)” on page 113</p> <p>“Managing Software With the Oracle Solaris Product Registry GUI (Task Map)” on page 397</p> <p>Chapter 5, “Setting Up Printers by Using LP Print Commands (Tasks),” in <i>System Administration Guide: Printing</i></p> <p>“Setting Up Terminals and Modems With Serial Ports Tool (Overview)” in <i>System Administration Guide: Advanced Administration</i></p>
Package and Patch Tools	<p>Starting with Oracle Solaris 10, the package and patch tools have been enhanced. Use the <code>pkgchk</code> command with the <code>-P</code> option instead of <code>grep pattern /var/sadm/install/contents</code>. The <code>-P</code> option enables you to use a partial path.</p>	<p>“Package and Patch Tool Enhancements” on page 388</p> <p>Chapter 23, “Managing Patches”</p>
Solaris Print Manager	<p>Expanded printer support in Solaris Print Manager includes the following features, which were introduced in Oracle Solaris 10:</p> <ul style="list-style-type: none"> ■ Never Print Banner option ■ Support for raster image processor (RIP) ■ Support for PostScript Printer Description (PPD) files: <ul style="list-style-type: none"> ■ The <code>-n</code> option of the <code>lpadmin</code> command enables you to specify a PPD file when creating or modifying printers. ■ The Use PPD Files option in Solaris Print Manager enables you to specify a PPD file when creating or modifying printers. ■ The <code>lpstat</code> command output displays the PPD file for a print queue that uses a PPD file. 	<p>“What’s New in Printing?” in <i>System Administration Guide: Printing</i></p>

Matrix of Oracle Solaris Management Tools and Supported Releases

This section provides information about tools that are primarily used to manage users, groups, clients, disks, printers, and serial ports.

The following table lists the various Oracle Solaris management GUI tools and whether they are currently supported.

TABLE 1-2 Matrix of Solaris Management Tool Support

	Solaris 9	Solaris 10
admintool	Supported	Not supported
Solstice AdminSuite 2.3	Not supported	Not supported
Solstice AdminSuite 3.0	Not supported	Not supported
Solaris Management Tools 1.0	Not supported	Not supported
Solaris Management Tools 2.0	Not supported	Not supported
Solaris Management Tools 2.1	Supported	Supported

If you need to perform administration tasks on a system that has a text-based terminal as the console, use Solaris Management Console commands instead. For more information, see [Table 1-5](#).

Feature Descriptions for Oracle Solaris Management Tools

This table describes the tools that are available in the Oracle Solaris release.

TABLE 1-3 Descriptions of Solaris Management Tools

Feature or Tool	Supported in Solaris Management Console 2.1?
Computers and Networks tool	Supported
Diskless Client support	A diskless client command-line interface is available
Disks tool	Supported
Enhanced Disk tool (Solaris Volume Manager)	Supported
Job Scheduler tool	Supported
Log Viewer tool	Supported

TABLE 1-3 Descriptions of Solaris Management Tools (Continued)

Feature or Tool	Supported in Solaris Management Console 2.1?
Mail Alias support	Supported
Mounts and Shares tool	Supported
Name Service support	For users, groups, and network information only
Performance tool	Supported
Printer support	Not Supported, but Solaris Print Manager is available as a separate tool
Projects tool	Supported
Role-based access control (RBAC) support	Supported
RBAC Tool	Supported
Serial Port tool	Supported
Software Package tool	Not supported
System Information tool	Supported
User/Group tool	Supported

Feature Descriptions for Solaris 9 Management Tools

This table describes the tools that are available in the Solaris 9 releases.

TABLE 1-4 Feature Descriptions for Solaris 9 Management Tools

Feature or Tool	Supported in admintool?	Supported in Solaris Management Console 2.1?
Computers and Networks tool	Not supported	Supported
Diskless Client support	Not supported	A diskless client command-line interface is available
Disks tool	Not supported	Supported
Enhanced Disk tool (Solaris Volume Manager)	Not supported	Supported
Job Scheduler tool	Not supported	Supported
Log Viewer tool	Not supported	Supported
Mail Alias support	Not supported	Supported
Mounts and Shares tool	Not supported	Supported

TABLE 1-4 Feature Descriptions for Solaris 9 Management Tools (Continued)

Feature or Tool	Supported in <code>admintool</code> ?	Supported in Solaris Management Console 2.1?
Name Service support	Not supported	For users, groups, and network information only
Performance tool	Not supported	Supported
Printer support	Supported	Not supported, but Solaris Print Manager is available as a separate tool
Projects tool	Not supported	Supported
RBAC support	Not supported	Supported
RBAC tool	Not supported	Supported
Serial Port tool	Supported	Supported
Software Package tool	Supported	Not supported
System Information tool	Not supported	Supported
User/Group tool	Supported	Supported

Availability of Solaris Management Commands

The following tables list commands that perform the same tasks as the Oracle Solaris management tools. For information on diskless client support, see [Chapter 7, “Managing Diskless Clients \(Tasks\)”](#).

Solaris 10 System Management Commands

This table describes the commands that provide the same functionality as the Oracle Solaris management tools. You must become superuser or assume an equivalent role to use these commands. Some of these commands are for the local system only. Others commands operate in a name service environment. See the appropriate man page and refer to the `-D` option.

TABLE 1-5 Descriptions for Solaris Management Commands

Command	Description	Man Page
<code>smc</code>	Starts the Solaris Management Console	smc(1M)
<code>smcron</code>	Manages crontab jobs	smcron(1M)

TABLE 1-5 Descriptions for Solaris Management Commands (Continued)

Command	Description	Man Page
<code>smdiskless</code>	Manages diskless client support	smdiskless(1M)
<code>smexec</code>	Manages entries in the <code>exec_attr</code> database	smexec(1M)
<code>smgroup</code>	Manages group entries	smgroup(1M)
<code>smlog</code>	Manages and views WBEM log files	smlog(1M)
<code>smmultiuser</code>	Manages bulk operations on multiple user accounts	smmultiuser(1M)
<code>smosservice</code>	Adds OS services and diskless client support	smosservice(1M)
<code>smprofile</code>	Manages profiles in the <code>prof_attr</code> and <code>exec_attr</code> databases	smprofile(1M)
<code>smrole</code>	Manages roles and users in role accounts	smrole(1M)
<code>smserialport</code>	Manages serial ports	smserialport(1M)
<code>smuser</code>	Manages user entries	smuser(1M)

This table describes the commands you can use to manage RBAC from the command line. You must be superuser or assume an equivalent role to use these commands. These commands cannot be used to manage RBAC information in a name service environment.

TABLE 1-6 RBAC Command Descriptions

Command	Description	References
<code>auths</code>	Displays authorizations granted to a user	auths(1)
<code>profiles</code>	Displays execution profiles for a user	profiles(1)
<code>roleadd</code>	Adds a new role to the system	roleadd(1M)
<code>roles</code>	Displays roles granted to a user	roles(1)

This table describes the commands you can use to manage users, groups, and RBAC features from the command line. You must be superuser or assume an equivalent role to use these commands. These commands cannot be used to manage user and group information in a name service environment.

TABLE 1-7 Oracle Solaris User and Group Command Descriptions

Command	Description	References
useradd, usermod, userdel	Adds, modifies, or removes a user	useradd(1M) , usermod(1M) , userdel(1M)
groupadd, groupmod, groupdel	Adds, modifies, or removes a group	groupadd(1M) , groupmod(1M) , groupdel(1M)

For More Information About Oracle Solaris Management Tools

This table identifies where to find more information about management tools in the Oracle Solaris release.

TABLE 1-8 For More Information About Solaris Management Tools

Tool	Availability	For More Information
Solaris Management Console 2.1 suite of tools	Solaris 9 and Oracle Solaris 10 releases	This guide and the console online help
Solaris Management Console 2.0 suite of tools	Solaris 8 1/01, 4/01, 7/01, 10/01, and 2/02 releases	Solaris Management Console online help
admintool	Solaris 9 and previous Solaris releases	admintool
AdminSuite 3.0	Solaris 8, Solaris 8 6/00, and Solaris 8 10/00 releases	<i>Solaris Easy Access Server 3.0 Installation Guide</i>
Diskless Client command-line interface	Solaris 8 1/01, 4/01, 7/01, 10/01, 2/02, Solaris 9 and Oracle Solaris 10	Chapter 7, “Managing Diskless Clients (Tasks)”

Working With the Solaris Management Console (Tasks)

This chapter describes the management tools that are used to perform system administration tasks. Topics include starting the Solaris Management Console (console), setting up role-based access control (RBAC) to use with the console, and working with the Solaris management tools in a name service environment.

For information about the procedures that are associated with performing system management tasks by using the Solaris Management Console, see the following task maps:

- [“Using the Solaris Management Tools With RBAC \(Task Map\)” on page 40](#)
- [“Using the Oracle Solaris Management Tools in a Name Service Environment \(Task Map\)” on page 45](#)

For information about troubleshooting Solaris Management Console problems, see [“Troubleshooting the Solaris Management Console” on page 52](#).

Solaris Management Console (Overview)

The following sections provide overview information about the Solaris Manager Console.

What Is the Solaris Management Console?

The Solaris Management Console is a container for GUI-based management tools that are stored in collections referred to as *toolboxes*.

The console includes a default toolbox with many basic management tools, including tools for managing the following:

- Users
- Projects
- cron jobs for mounting and sharing file systems

- cron jobs for managing disks and serial ports

For a brief description of each Solaris management tool, see [Table 2-1](#).

You can add tools to the existing toolbox, or you can create new toolboxes.

The Solaris Management Console has three primary components:

- **The Solaris Management Console client**
Called the *console*, this component is the visible interface and contains the GUI tools that are used to perform management tasks.
- **The Solaris Management Console server**
This component is located either on the same system as the console or it is remote. This component provides all of the *back-end* functionality that enables management through the console.
- **The Solaris Management Console toolbox editor**
This application, which appears similar to the console, is used to add or modify toolboxes, to add tools to a toolbox, or to extend the scope of a toolbox. For example, you could add a toolbox to manage a name service domain.

The default toolbox is visible when you start the console.

Solaris Management Console Tools

This table describes the tools that are included in the default Solaris Management Console toolbox. Cross-references to background information for each tool are provided.

TABLE 2-1 Solaris Management Console Tool Suite

Category	Tool	Description	For More Information
System Status	System Information	Monitors and manages system information such as date, time, and time zone	Chapter 5, “Displaying and Changing System Information (Tasks),” in <i>System Administration Guide: Advanced Administration</i>
	Log Viewer	Monitors and manages the Solaris Management Console tools log and system logs	Chapter 14, “Troubleshooting Software Problems (Overview),” in <i>System Administration Guide: Advanced Administration</i>
	Processes	Monitors and manages system processes	“Processes and System Performance” in <i>System Administration Guide: Advanced Administration</i>

TABLE 2-1 Solaris Management Console Tool Suite			(Continued)
Category	Tool	Description	For More Information
	Performance	Monitors system performance	Chapter 11, “Managing System Performance (Overview);” in <i>System Administration Guide: Advanced Administration</i>
System Configuration	Users	Manages users, rights, roles, groups, and mailing lists	“What Are User Accounts and Groups?” on page 86 and “Role-Based Access Control (Overview)” in <i>System Administration Guide: Security Services</i>
	Projects	Creates and manages entries in the <code>/etc/project</code> database	Chapter 2, “Projects and Tasks (Overview);” in <i>System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones</i>
	Computers and Networks	Creates and monitors computer and network information	Solaris Management Console online help
Services	Scheduled Jobs	Creates and manages scheduled cron jobs	“Ways to Automatically Execute System Tasks” in <i>System Administration Guide: Advanced Administration</i>
Storage	Mounts and Shares	Mounts and shares file systems	“Mounting and Unmounting Oracle Solaris File Systems” in <i>System Administration Guide: Devices and File Systems</i>
	Disks	Creates and manages disk partitions	Chapter 10, “Managing Disks (Overview);” in <i>System Administration Guide: Devices and File Systems</i>
	Enhanced Storage	Creates and manages volumes, hot spare pools, state database replicas, and disk sets	<i>Solaris Volume Manager Administration Guide</i>
Devices and Hardware	Serial Ports	Sets up terminals and modems	Chapter 1, “Managing Terminals and Modems (Overview);” in <i>System Administration Guide: Advanced Administration</i>

Context-sensitive help is available after you start a tool. For more in-depth online information than the context help provides, see the expanded help topics. You can access these topics from the console Help menu.

Why Use the Solaris Management Console?

The console provides a set of tools that have many benefits for administrators.

The console does the following:

- **Supports all experience levels**

Inexperienced administrators can complete tasks by using the GUI, which includes dialog boxes, wizards, and context help. Experienced administrators find that the console provides a convenient, secure alternative to using a text editor to manage hundreds of configuration parameters spread across dozens, or hundreds, of systems.

- **Controls user access to the system**

Although any user can access the console by default, only superuser can make changes to the initial configuration. As described in [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services*, it is possible to create special user accounts called, *roles*, which can be assigned to users, typically administrators, who are permitted to make specific system changes.

The key benefit of RBAC is that roles can be limited, so that users have access to only those tasks that are necessary for doing their jobs. RBAC is *not* required for using the Solaris management tools. You can run all of the tools as superuser, without making any changes.

- **Provides a command-line interface**

If preferred, administrators can operate the Solaris management tools through a command-line interface (CLI). Some commands are written specifically to mimic the GUI tool functions, such as the commands for managing users. These commands are listed in [Table 1–5](#), which includes the names and brief descriptions of each command. There is also a man page for each command.

For Solaris management tools that have no special commands, such as the Mounts and Shares tool, use the standard UNIX commands.

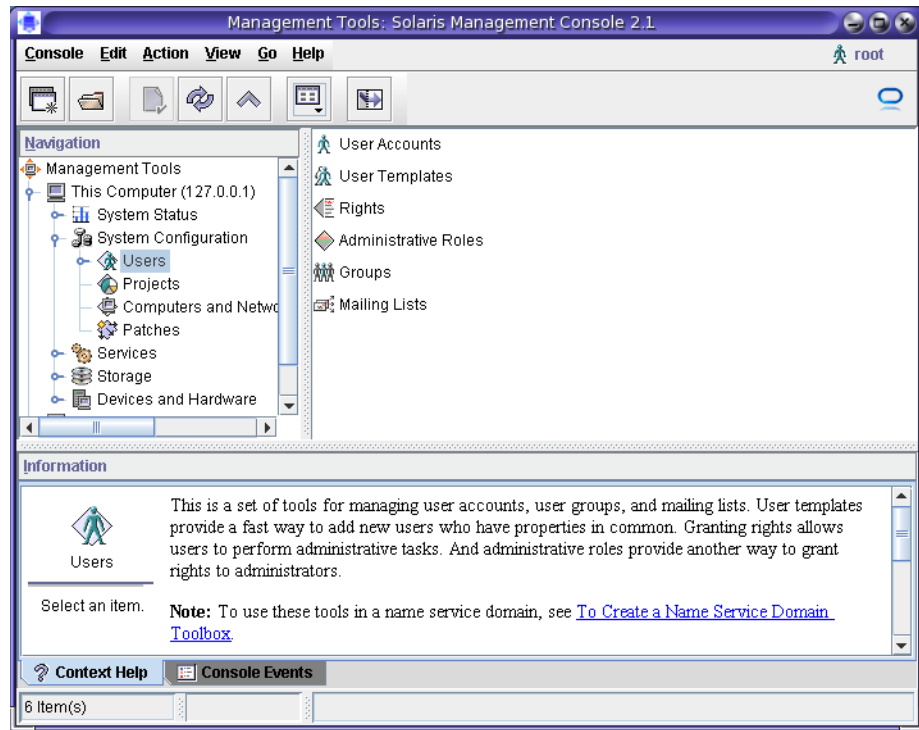
For in-depth information about how RBAC works, its benefits, and how to apply those benefits to your site, see [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services*.

To learn more about using RBAC with the Oracle Solaris management tools, see [“Using the Solaris Management Tools With RBAC \(Task Map\)”](#) on page 40.

Organization of the Solaris Management Console

In the following figure, the console is shown with the Users tool open.

FIGURE 2-1 Solaris Management Console Users Tool



The main part of the console consists of three panes:

- **Navigation pane** (left) – For accessing tools (or sets of tools), folders, or other toolboxes. Icons in the navigation pane are called *nodes* and are expandable, if they are folders or toolboxes.
- **View pane** (right) – For viewing information related to the node selected in the navigation pane. The view pane shows either the contents of the selected folder, subordinate tools, or the data that is associated with the selected tool.
- **Information pane** (bottom) – For displaying context-sensitive help or console events.

Changing the Solaris Management Console Window

The layout of the console window is highly configurable. You can use the following features to change the console window layout:

- **View menu** – Use the Show option in the View menu to hide or display the optional bars and panes. The other options in the View menu control the display of nodes in the view pane.
- **Console menu** – Use the Preferences option to set the following: the initial toolbox, the orientation of panes, clicking or double-clicking for selection, text or icons in the tool bar, fonts, default tool loading, authentication prompts, and advanced logins.
- **Context Help or Console Events toggles** – Use the icons at the bottom of the information pane to toggle between the display of context-sensitive help and console events.

Solaris Management Console Documentation

The main source of documentation for using the console and its tools is the online help system. The following two forms of online help are available:

- **Context-sensitive help responds to your use of the console tools.**
Clicking the cursor on tabs, entry fields, radio buttons, and so forth, causes the appropriate help to be displayed in the Information pane. You can close or reopen the Information pane by clicking the question mark button on dialog boxes and wizards.
- **Expanded help topics are available from the Help menu or by clicking cross reference links in some of the context-sensitive help.**
These topics, which are displayed in a separate viewer, contain more in-depth information than is provided by the context help. Topics include: overviews of each tool, explanations of how each tool works, files that are used by a specific tool, and troubleshooting information.

For a brief overview of each tool, refer to [Table 2-1](#).

How Much Role-Based Access Control?

As described in “[Why Use the Solaris Management Console?](#)” on page 34, a major advantage of using the Solaris management tools is the ability to use Role-Based Access Control (RBAC). RBAC provides administrators with access to just the tools and commands they need to perform their jobs.

Depending on your security needs, you can use varying degrees of RBAC.

RBAC Approach	Description	For More Information
No RBAC	Enables you to perform all tasks as superuser. You can log in as yourself. When you select a Solaris management tool, you specify root as the user and the root password.	“How to Become Superuser (root) or Assume a Role” on page 38
root as a role	Eliminates anonymous root logins and prevents users from logging in as root. This approach requires users to log in as themselves before they assume the root role. Note that you can apply this approach whether or not you are using other roles.	“How to Plan Your RBAC Implementation” in <i>System Administration Guide: Security Services</i>
Single role only	Uses the Primary Administrator role, which is essentially equivalent to having root access.	“Creating the Primary Administrator Role” on page 41
Suggested roles	Uses three roles that are easily configured: Primary Administrator, System Administrator, and Operator. These roles are appropriate for organizations with administrators at different levels of responsibility whose job capabilities fit the suggested roles.	“Role-Based Access Control (Overview)” in <i>System Administration Guide: Security Services</i>
Custom roles	You can add your own roles, depending on your organization's security needs.	“Managing RBAC” in <i>System Administration Guide: Security Services</i> and “How to Plan Your RBAC Implementation” in <i>System Administration Guide: Security Services</i>

Becoming Superuser (root) or Assuming a Role

Most administration tasks, such as adding users or managing file systems require that you first log in as root (UID=0) or assume a role, if you are using RBAC. The root account, also known as the *superuser* account, is used to make system changes and can override user file protection in emergency situations.

The superuser account and roles should be used only to perform administrative tasks to prevent indiscriminate changes to the system. The security problem that is associated with the superuser account is that this user has complete access to the system, even when performing minor tasks.

In a non-RBAC environment, you can either log in to the system as superuser or use the `su` command to change to the superuser account. If RBAC is implemented, you can assume roles through the console or use `su` and specify a role.

When you use the console to perform administration tasks, you can do one of the following:

- Log in to the console as yourself and then supply the root user name and password
- Log in to the console as yourself and then assume a role

A major benefit of RBAC is that roles can be created to give limited access to specific functions only. If you are using RBAC, you can run restricted applications by assuming a role rather than by becoming superuser.

For step-by-step instructions on creating the Primary Administrator role, see [“How to Create the First Role \(Primary Administrator\)”](#) on page 42. For an overview of RBAC, see [Chapter 9, “Using Role-Based Access Control \(Tasks\)”](#) in *System Administration Guide: Security Services*.

▼ How to Become Superuser (root) or Assume a Role

Become superuser or assume a role by using one of the following methods. Each method requires that you know either the superuser password or the role password.

1 Become superuser by selecting one of the following methods:

- **Log in as a user, then do the following:**
 - a. **Start the Solaris Management Console.**
 - b. **Select a Solaris management tool.**
 - c. **Log in as root.**

This method enables to you perform any management task from the console.

For information on starting the Solaris Management Console, see [“How to Start the Solaris Management Console in a Name Service Environment”](#) on page 51.

- **Log in as superuser on the system console.**

```
hostname console: root
Password: root-password
#
```

The pound sign (#) is the shell prompt for the superuser account.

This method provides complete access to all of the system commands and tools.

- **Log in as a user, then change to the superuser account by using the su command at the command line.**

```
% su
Password: root-password
#
```

This method provides complete access to all of the system commands and tools.

- **Log in remotely as superuser.**

This method is not enabled by default. You must modify the `/etc/default/login` file to remotely log in as superuser on the system console. For information on modifying this file, see [Chapter 3, “Controlling Access to Systems \(Tasks\)”](#) in *System Administration Guide: Security Services*.

This method provides complete access to all system commands and tools.

2 Assume a role.

Select one of the following methods:

- **Log in as user, then change to a role by using the su command at the command line.**

```
% su role
Password: role-password
$
```

This method provides access to all of the commands and tools that the role has access to.

- **Log in as a user, then do the following:**

a. **Start the Solaris Management Console.**

b. **Select a Solaris management tool.**

c. **Assume a role.**

For information on starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role”](#) on page 44.

This method provides access to all of the Solaris management tools that the role has access to.

Using the Solaris Management Tools With RBAC (Task Map)

This task map describes the tasks you will need to perform, if you want to use the RBAC security features to perform administration tasks, rather than use the superuser account.

Note – The information in this chapter describes how to use the console with RBAC. RBAC overview and task information is included to show how to initially set up RBAC with the console.

For detailed information about RBAC and how to use it with other applications, see [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services*.

Task	Description	For Instructions
1. Start the console.	If your user account is already set up, start the console as yourself. Then, log in to the console as root. If you do not have a user account set up, become superuser first, and then start the console.	“How to Start the Console as Superuser or as a Role” on page 44
2. Add a user account for yourself.	Add a user account for yourself, if you do not have an account already.	Solaris Management Console online help “If You Are the First to Log In to the Console” on page 41
3. Create the Primary Administrator role	Create the Primary Administrator role. Then, add yourself to this role.	“How to Create the First Role (Primary Administrator)” on page 42
4. Assume the Primary Administrator role.	Assume the Primary Administrator role after you have created this role.	“How to Assume the Primary Administrator Role” on page 43
5. (Optional) Make root a role.	Make root a role and add yourself to the root role, so that no other user can use the su command to become root.	“How to Plan Your RBAC Implementation” in <i>System Administration Guide: Security Services</i>
6. (Optional) Create other administrative roles.	Create other administrative roles and grant the appropriate rights to each role. Then, add the appropriate users to each role.	Chapter 9, “Using Role-Based Access Control (Tasks)” , in <i>System Administration Guide: Security Services</i>

The following sections provide overview information and step-by-step instructions for using the Solaris Management Console and the RBAC security features.

If You Are the First to Log In to the Console

If you are the first administrator to log in to the console, start the console as a user (yourself). Then, log in as superuser. This method gives you complete access to all of the console tools.

Here are the general steps to follow, depending on whether you are using RBAC:

- **Without RBAC** – If you choose not to use RBAC, continue working as superuser. All other administrators will also need root access to perform their jobs.
- **With RBAC** – You will need to do the following:
 - If you do not already have an account, set up your user account.
 - Create the role called Primary Administrator.
 - Assign the Primary Administrator right to the role that you are creating.
 - Assign your user account to this role.

For step-by-step instructions on creating the Primary Administrator role, see [“How to Create the First Role \(Primary Administrator\)”](#) on page 42.

For an overview of how RBAC works, see [Chapter 9, “Using Role-Based Access Control \(Tasks\)”](#) in *System Administration Guide: Security Services*.

Creating the Primary Administrator Role

An *administrator role* is a special user account. Users who assume a role are permitted to perform a predefined set of administrative tasks.

The Primary Administrator role is permitted to perform all administrative functions, similar to superuser.

If you are superuser, or a user who is assuming the Primary Administrator role, you can define which tasks other administrators are permitted to perform. With the help of the Add Administrative Role wizard, you can create a role, grant rights to the role, and then specify which users are permitted to assume that role. A *right* is a named collection of commands, or authorizations, for using specific applications. A right enables you to perform specific functions within an application. The use of rights can be granted or denied by an administrator.

The following table describes the information that you are prompted for when you create the Primary Administrator role.

TABLE 2-2 Field Descriptions for Adding a Role by Using the Solaris Management Console

Field name	Description
Role name	Selects the name an administrator uses to log in to a specific role.
Full name	Provides a full, descriptive name of this role. (Optional)

TABLE 2-2 Field Descriptions for Adding a Role by Using the Solaris Management Console
(Continued)

Field name	Description
Description	Provides further description of this role.
Role ID number	Selects the identification number assigned to this role. This number is the same as the set of identifiers for UIDs.
Role shell	Selects the shell that runs when a user logs in to a terminal or console window and assumes a role in that window.
Create a role mailing list	Creates a mailing list with the same name as the role, if checked. You can use this list to send email to everyone assigned to the role.
Role password and confirm Password	Sets and confirms the role password.
Available rights and granted Rights	Assigns rights to this role by choosing from the list of Available Rights and adding them to the list of Granted Rights.
Select a home directory	Selects the home directory server where this role's private files will be stored.
Assign users to this role	Adds specific users to the role so that they can assume the role to perform specific tasks.

For detailed information about RBAC and instructions on how to use roles to create a more secure environment, see [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services*.

▼ How to Create the First Role (Primary Administrator)

This procedure describes how to create the Primary Administrator role and then assign it to your user account. This procedure assumes that your user account is already created.

1 Start the console as yourself.

```
% /usr/sadm/bin/smc &
```

For additional information about starting the console, see [“How to Start the Console as Superuser or as a Role”](#) on page 44.

The console online help provides more information about creating a user account for yourself.

2 Click the This Computer icon in the Navigation pane.

3 Click System Configuration ->Users -> Administrative Roles.

4 Click Action ->Add Administrative Role.

The Add Administrative Role wizard opens.

5 Create the Primary Administrator role with the Administrative Role wizard by following these steps:

- a. Identify the role name, which includes the full role name, description, role ID number, role shell, and whether you want to create a role mailing list. Click Next.
- b. Set and confirm the role password, then click Next.
- c. Select the Primary Administrator right from the Available Rights column, add it to Granted Rights column.
- d. Click Next.
- e. Select the home directory for the role, then click Next.
- f. Assign yourself to the list of users who can assume the role, then click Next.

If necessary, see [Table 2-2](#) for a description of the role fields.

6 Click Finish.

▼ How to Assume the Primary Administrator Role

After you have created the Primary Administrator role, you will need to log in to the console as yourself, and then assume the Primary Administrator role. When you assume a role, you take on all of the attributes of that role, including the rights. At the same time, you relinquish all of your own user properties.

1 Start the console.

```
% /usr/sadm/bin/smc &
```

For information about starting the console, see [“How to Start the Console as Superuser or as a Role” on page 44](#).

2 Log in with your user name and password.

A list shows which roles you are permitted to assume.

3 Log in to the Primary Administrator role and provide the role password.

Starting the Solaris Management Console

The following procedure describes how to start the console and gain access to the Solaris management tools.

For instructions on what to do if you are the first user to log in to the console, see [“If You Are the First to Log In to the Console” on page 41](#).

▼ How to Start the Console as Superuser or as a Role

If you start the console as a user with your own user account, you have limited access to the Solaris management tools. For greater access, you can log in as yourself, and then log in as one of the roles that you are permitted to assume. If you are permitted to assume the role of Primary Administrator, you have access to all of the Solaris management tools. This role is equivalent to superuser.

- 1 **Verify that you are in a window environment, such as the GNOME environment.**

- 2 **Start the console.**

```
% /usr/sadm/bin/smc &
```

It might take a minute or two for the console to come up the first time.

The Solaris Management Console window is displayed.

Note – Open a console in your window environment to display the Solaris Management Console startup messages. Do not attempt to start the Solaris Management Console server manually before starting the Solaris Management Console application. The server starts automatically when you start the Solaris Management Console. For information on troubleshooting console problems, see [“Troubleshooting the Solaris Management Console” on page 52](#).

- 3 **Under the Management Tools icon in the Navigation pane, double-click the This Computer icon.**

A list of categories is displayed.

- 4 **(Optional) Select the appropriate toolbox.**

If you want to use a toolbox other than the default toolbox, select that toolbox from the Navigation pane. Or, select Open Toolbox from the console menu and load the toolbox that you want.

For information about using different toolboxes, see [“How to Create a Toolbox for a Specific Environment” on page 48](#).

- 5 **To access a particular tool, double-click the category icon.**

Use the online help to identify how to perform a specific task.

6 Double-click the tool icon.

A pop-up Log-In window is displayed.

7 Decide if you want to use the tool as superuser or as a role.

- **If you are logging in as superuser, enter the root password.**
- **If you are logging in as yourself, backspace over the root user name, then supply your user ID and user password.**

A list of roles you can assume is displayed.

8 Select the Primary Administrator role or an equivalent role, then supply the role password.

For step-by-step instructions on creating the Primary Administrator role, see [“How to Create the First Role \(Primary Administrator\)”](#) on page 42.

The main tool menu is displayed.

Using the Oracle Solaris Management Tools in a Name Service Environment (Task Map)

By default, the Oracle Solaris management tools are set up to operate in a local environment. For example, the Mounts and Shares tool enables you to mount and share directories on specific systems, but not in an NIS or NIS+ environment. However, you can manage information with the Users and Computers and Networks tools in a name service environment.

To work with a console tool in a name service environment, you need to create a name service toolbox, and then add the tool to that toolbox.

Task	Description	For Instructions
1. Verify prerequisites.	Verify you have completed the prerequisites before attempting to use the console in a name service environment.	“Prerequisites for Using the Solaris Management Console in a Name Service Environment” on page 47
2. Create a toolbox for the name service.	Use the New Toolbox wizard to create a toolbox for your name service tools.	“How to Create a Toolbox for a Specific Environment” on page 48
3. Add a tool to the name service toolbox.	Add the Users tool, or any other name service tool, to your name service toolbox.	“How to Add a Tool to a Toolbox” on page 50

Task	Description	For Instructions
4. Select the toolbox that was just created.	Select the toolbox you just created to manage name service information.	“How to Start the Solaris Management Console in a Name Service Environment” on page 51

RBAC Security Files

The RBAC security files that work with the Solaris Management Console are created when you upgrade to or install at least the Solaris 9 release. If you do not install the Solaris Management Console packages, the RBAC security files are installed without the necessary data for using RBAC. For information on the Solaris Management Console packages, see [“Troubleshooting the Solaris Management Console” on page 52](#).

The RBAC security files if you are running at least the Solaris 9 release are included in your name service so that you can use the Solaris Management Console tools in a name service environment.

The security files on a local server are populated into a name service environment as part of a standard upgrade by the `ypmake`, `nispopulate`, or equivalent LDAP commands.

The following name services are supported:

- NIS
- NIS+
- LDAP
- files

The RBAC security files are created when you upgrade to or install Oracle Solaris 10.

This table briefly describes the predefined security files that are installed on a system that is running the Oracle Solaris release.

TABLE 2-3 RBAC Security Files

Local File Name	Table or Map Name	Description
<code>/etc/user_attr</code>	<code>user_attr</code>	Associates users and roles with authorizations and rights profiles
<code>/etc/security/auth_attr</code>	<code>auth_attr</code>	Defines authorizations and their attributes and identifies associated help files
<code>/etc/security/prof_attr</code>	<code>prof_attr</code>	Defines rights profiles, lists the rights profiles assigned to the authorizations, and identifies associated help files

TABLE 2-3 RBAC Security Files (Continued)

Local File Name	Table or Map Name	Description
/etc/security/exec_attr	exec_attr	Defines the privileged operations assigned to a rights profile

For unusual upgrade cases, you might have to use the `smat rpop` command to populate RBAC security files in the following instances:

- When creating or modifying rights profiles
- When you need to include users and roles by customizing the `usr_attr` file

For more information, see “[Role-Based Access Control \(Overview\)](#)” in *System Administration Guide: Security Services*.

Prerequisites for Using the Solaris Management Console in a Name Service Environment

The following table identifies what you need to do before you can use the Solaris Management Console in a name service environment.

Prerequisite	For More Information
Install the Oracle Solaris 10 release.	<i>Oracle Solaris 10 8/11 Installation Guide: Basic Installations</i>
Set up your name service environment.	<i>System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)</i>
Select your management scope.	“Management Scope” on page 47
Make sure your <code>/etc/nsswitch.conf</code> file is configured, so that you can access your name service data.	“ <code>/etc/nsswitch.conf</code> File” on page 48

Management Scope

The Solaris Management Console uses the term *management scope* to refer to the name service environment that you want to use with the selected management tool. The management scope choices for the Users tool and the Computers and Networks tool are LDAP, NIS, NIS+, or files.

The management scope that you select during a console session should correspond to the primary name service that is identified in the `/etc/nsswitch.conf` file.

/etc/nsswitch.conf File

The `/etc/nsswitch.conf` file on each system specifies the policy for name service lookups (where data is read from) on that system.

Note – You must make sure that the name service accessed from the console, which you specify through the console Toolbox Editor, appears in the search path of the `/etc/nsswitch.conf` file. If the specified name service does not appear there, the tools might behave in unexpected ways, resulting in errors or warnings.

When you use the Solaris management tools in a name service environment, you might impact many users with a single operation. For example, if you delete a user in the NIS or NIS+ name service, that user is deleted on all systems that are using NIS or NIS+.

If different systems in your network have different `/etc/nsswitch.conf` configurations, unexpected results might occur. So, all systems to be managed with the Solaris management tools should have a consistent name service configuration.

▼ **How to Create a Toolbox for a Specific Environment**

Applications for administering the Oracle Solaris operating system are called tools. Those tools are stored in collections referred to as *toolboxes*. A toolbox can be located on a local server where the console is located or on a remote machine.

Use the Toolbox Editor to do the following:

- Add a new toolbox
- Add tools to an existing toolbox
- Change the scope of a toolbox

For example, use this tool to change the domain from local files to a name service.

Note – You can start the Toolbox Editor as a regular user. However, if you plan to make changes and save them to the default console toolbox, `/var/sadm/smc/toolboxes`, you must start the Toolbox Editor as `root`.

- 1 Start the Toolbox Editor.**
`# /usr/sadm/bin/smc edit &`
- 2 Select Open from the Toolbox menu.**
- 3 In the Toolboxes window, select This Computer.**

- 4 Click Open.**

The This Computer toolbox opens.
- 5 In the Navigation pane, select the This Computer icon again.**
- 6 From the Action menu, select Add Folder.**
- 7 Use the Folder wizard to add a new toolbox for your name service environment.**
 - a. Name and Description – Provide a name in the Full Name window, then click Next.**

For example, for the NIS environment, provide “NIS tools”.
 - b. Provide a description in the Description window, then click Next.**

For example, “tools for NIS environment” is an appropriate description.
 - c. Icons – Use the default value for the Icons, then click Next.**
 - d. Management Scope – Select Override.**
 - e. Under the Management Scope pull-down menu, Select your name service u.**
 - f. Add the name service master name in the Server field, if necessary.**
 - g. In the Domain field, add the domain that is managed by the server.**
 - h. Click Finish.**

The new toolbox is displayed in the left Navigation pane.
- 8 Select the new toolbox icon, then select Save As from the Toolbox menu.**
- 9 In the Local Toolbox Filename dialog, enter the toolbox path name.**

Use the .tbx suffix.

```
/var/sadm/smc/toolboxes/this_computer/toolbox-name.tbx
```
- 10 Click Save.**

The new toolbox is displayed in the Navigation pane in the console window.

See Also After you have created a name service toolbox, you can put a name service tool into it. For more information, see [“How to Add a Tool to a Toolbox” on page 50](#).

▼ How to Add a Tool to a Toolbox

In addition to the default tools that ship with the console, additional tools can be launched from the console. As these tools become available, you can add one or more tools to an existing toolbox.

You can also create a new toolbox for either local management or network management. Then, you can add tools to the new toolbox.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.

2 Start the Toolbox Editor, if necessary.

```
# /usr/sadm/bin/smc edit &
```

3 Select the toolbox.

If you want to work in a name service, select the toolbox that you just created in the Toolbox Editor. For more information, see [“How to Create a Toolbox for a Specific Environment”](#) on page 48.

4 From the Action menu, select Add Tool.

5 Use the Add Tool wizard to add the new tool.

a. **Server Selection** – Add the name service master in the Server window. Click Next.

b. **Tools Selection** – Select the tool you want to add from the Tools window. Click Next.

If this toolbox is a name service toolbox, choose a tool that you want to work with in the name service environment. For example, choose the Users tool.

c. **Name and Description** – Accept the default values, then click Next.

d. **Icons** – Accept the default values, unless you have created custom icons. Click Next.

e. **Management Scope** – Accept the default value, “Inherit from Parent.” Click Next.

f. **Tool Loading** – Accept the default value, “Load tool when selected.” Click Finish.

6 To save the updated toolbox, Select Save.

The Local Toolbox window is displayed.

▼ How to Start the Solaris Management Console in a Name Service Environment

After you have created a name service toolbox and added tools to it, you can start the Solaris Management Console and open that toolbox to manage a name service environment.

Before You Begin Verify that the following prerequisites are met:

- Ensure that the system you are logged in to is configured to work in a name service environment.
- Verify that the `/etc/nsswitch.conf` file is configured to match your name service environment.

1 Start the Solaris Management Console.

For more information, see [“How to Start the Console as Superuser or as a Role”](#) on page 44.

2 Select the toolbox that you created for the name service

The toolbox is displayed in the Navigation pane.

For information about creating a toolbox for a name service, see [“How to Create a Toolbox for a Specific Environment”](#) on page 48.

Adding Tools to the Solaris Management Console

You can add legacy tools or unbundled tools to the console. If you want to add authentication to these tools, see [“Managing RBAC”](#) in *System Administration Guide: Security Services*.

▼ How to Add a Legacy Tool to a Toolbox

A legacy tool is any application that was not designed specifically as a Solaris management tool. Each tool you add to a toolbox can then be launched from the Solaris Management Console.

You can add the following types of legacy tool applications to a console toolbox:

- X applications
- Command-line interface (CLI) applications
- HTML applications

1 Become superuser or assume an equivalent role.

2 Start the Solaris Management Console Toolbox Editor, if necessary.

```
# /usr/sadm/bin/smc edit &
```

- 3 Open the toolbox to which you want to add the legacy application.**
The toolbox selected is opened in the Toolbox Editor.
- 4 Select the node in the toolbox to which you want to add the legacy application.**
A legacy application can be added to the top node of a toolbox or to another folder.
- 5 Click Action -> Add Legacy Application.**
The General panel of the Legacy Application Wizard is displayed.
- 6 Follow the instructions in the wizard.**
- 7 Save the toolbox in the Toolbox Editor.**

▼ How to Install an Unbundled Tool

If you want to add a new tool package that can be launched from the Solaris Management Console, use the following procedure.

- 1 Become superuser or assume an equivalent role.**
- 2 Install the new tool package.**

```
# pkgadd ABCDtool
```
- 3 Restart the console, so that it recognizes the new tool.**
 - a. Stop the console server.**

```
# /etc/init.d/init.wbem stop
```
 - b. Start the console server.**

```
# /etc/init.d/init.wbem start
```
- 4 To verify that the new tool is displayed, start the console.**
For more information, see [“How to Start the Console as Superuser or as a Role”](#) on page 44.

Troubleshooting the Solaris Management Console

Before following this troubleshooting procedure, make sure that the following packages are installed:

- SUNWmc – Solaris Management Console 2.1 (Server Components)
- SUNWmcc – Solaris Management Console 2.1 (Client Components)
- SUNWmc.com – Solaris Management Console 2.1 (Common Components)

- SUNWmcdev – Solaris Management Console 2.1 (Development Kit)
- SUNWmcex – Solaris Management Console 2.1 (Examples)
- SUNWwbmc – Solaris Management Console 2.1 (WBEM Components)

These packages provide the basic Solaris Management Console launcher. Note that you must install the SUNWCprog cluster to use the Solaris Management Console and all of its tools.

▼ How to Troubleshoot the Solaris Management Console

The client and the server are started automatically when you start the Solaris Management Console.

If the console is visible and you are having trouble running the tools, the server might not be running, or the server might be in a problem state that can be resolved by stopping and restarting it.

1 Become superuser or assume an equivalent role.

2 Determine whether the console server is running.

```
# /etc/init.d/init.wbem status
```

If the console server is running, you should see a message similar the following:

```
SMC server version 2.1.0 running on port 898.
```

3 If the console server is not running, start it.

```
# /etc/init.d/init.wbem start
```

After a short time, you should see a message similar to the following:

```
SMC server is ready.
```

4 If the server is running, and you are still having problems, do the following:

a. Stop the console server.

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

You should see a message similar to the following:

```
Shutting down SMC server on port 898.
```

b. Start the console server.

```
# /etc/init.d/init.wbem start
```


Working With the Oracle Java Web Console (Tasks)

This chapter describes the Oracle Java Web Console, which is used to administer web-based Sun system management applications that are installed and registered on your system.

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

- “What's New in Administering the Oracle Java Web Console?” on page 55
- “Oracle Java Web Console (Overview)” on page 56
- “Getting Started With the Oracle Java Web Console” on page 59
- “Managing the Console Service” on page 62
- “Configuring the Oracle Java Web Console” on page 64
- “Troubleshooting the Oracle Java Web Console Software” on page 71
- “Oracle Java Web Console Reference Information” on page 78

For information about the procedures that are associated with using the Oracle Java Web Console, see “Getting Started With the Oracle Java Web Console (Task Map)” on page 58 and “Troubleshooting the Oracle Java Web Console Software (Task Map)” on page 70.

What's New in Administering the Oracle Java Web Console?

This section includes features that are new in this Oracle Solaris release. For a complete listing of new features and a description of Oracle Solaris releases, see [Oracle Solaris 10 8/11 What's New](#).

Oracle Java Web Console Server Management

Solaris 10 11/06: The Oracle Java Web Console server is managed as a service by the Service Management Facility (SMF). For more information about SMF, see [Chapter 18, “Managing Services \(Overview\)”](#).

Applications That Are Available to the Oracle Java Web Console

Solaris 10 6/06: The Oracle Solaris ZFS web-based management tool is available in the Oracle Java Web Console. This tool enables you to perform much of the administration tasks that you can perform with the command-line interface (CLI). These capabilities include setting parameters, viewing the various pools and file systems, and making updates to them.

The following are examples of typical procedures that you might perform with the tool:

- Create a new storage pool.
- Add capacity to an existing pool.
- Move (export) a storage pool to another system.
- Import a previously exported storage pool, to make it available on another system.
- View tables of information about storage pools.
- Create a file system.
- Create a zvol (virtual volume).
- Take a snapshot of a file system or a zvol volume.
- Roll back a file system to a previous snapshot.

For more information about using the Oracle Solaris ZFS web-based management tool, see [Oracle Solaris ZFS Administration Guide](#).

Note – The Java Enterprise System software includes several management applications that run in the Oracle Java Web Console.

Oracle Java Web Console (Overview)

The Oracle Java Web Console provides a common location for users to access web-based system management applications. You access the web console by logging in through a secure https port with one of several supported web browsers. The single entry point that the web console provides eliminates the need to learn URLs for multiple applications. In addition, the single entry point provides user authentication and authorization for all applications that are registered with the web console.

All web console-based applications conform to the same user interface guidelines, which enhances ease of use. The web console also provides auditing of user sessions and logging service for all registered applications.

What Is the Oracle Java Web Console?

The Oracle Java Web Console is a web page where you can find the Sun system management web-based applications that are installed and registered on your system. Registration is automatically a part of an application's installation process. Thus, registration requires no administrator intervention.

The Oracle Java Web Console provides the following:

- **A single point of entry for login and the launching of browser-based system management applications**

The console provides a central location from which you can start browser-based management applications simply by clicking the application names. No compatibility exists between the Oracle Java Web Console and the Solaris Management Console. The Oracle Java Web Console is a web application that you access through a browser, and Solaris Management Console is a Java application that you start from a command line. Because the consoles are completely independent, you can run both consoles on the same system at the same time.

- **Single sign-on through a secure https port**

Single sign-on in this context means that you do not have to authenticate yourself to each management application after you authenticate yourself to the web console. You enter your user name and password just once per console session.

- **Dynamically organized and aggregated applications**

Applications are installed and displayed on the console launch page under the category of management tasks that is most applicable.

Categories include the following:

- Systems
- Storage
- Services
- Desktop applications
- Other

- **A common look and feel**

All web console applications use the same user interface (UI) components and behavior, thereby reducing the learning curve for administrators.

- **Standard, extensible authentication, authorization, and auditing mechanisms**

The Oracle Java Web Console supports Pluggable Authentication Module (PAM), role-based access control (RBAC) roles, and Basic Security Module (BSM) auditing.

Oracle Java Web Console Management Commands

The Oracle Java Web Console includes the following management commands:

- `smcwebserver` – This command starts and stops the console's web server.
- `wcadmin` – **Starting with the Solaris 10 11/06 release**, this command is used to configure the console, and to register and deploy console applications. For more information, see the [wcadmin\(1M\)](#) man page.
- `smreg` – **In the Oracle Solaris 10, Solaris 10 1/06, and Solaris 10 6/06 OS**, this command is used to register all console applications.

Starting with the Solaris 10 11/06 release, use this command only to register legacy applications that were created for a version of the console that is not at least Oracle Java Web Console 3.0.

The commands are used to perform various tasks that this chapter describes.

For more information about each command, see the [smcwebserver\(1M\)](#), [wcadmin\(1M\)](#), and the [smreg\(1M\)](#) man pages.

Supported Web Browsers

The Oracle Java Web Console can be used in any of the following browsers while running Oracle Solaris:

- Mozilla (at least Version, 1.4)
- Netscape (at least Version, 6.2)
- Firefox (at least Version, 1.0)

Getting Started With the Oracle Java Web Console (Task Map)

Task	Description	For Instructions
Start applications from the Oracle Java Web Console's launch page.	The Oracle Java Web Console's launch page lists all of the registered system management applications that you have permission to use. You connect to a specific application by clicking its application name.	“How to Start Applications From the Oracle Java Web Console's Launch Page” on page 60

Task	Description	For Instructions
Start, stop, enable, and disable the console server.	You can manage the web server that is used to run the console and the registered applications.	<p>“How to Start the Console Service” on page 62</p> <p>“How to Enable the Console Service to Run at System Start” on page 62</p> <p>“How to Stop the Console Service” on page 63</p> <p>“How to Disable the Console Service” on page 63</p>
Change the Oracle Java Web Console's properties.	<p>You should not have to change any of the web console's default properties.</p> <p>Properties that you might choose to change include the following:</p> <ul style="list-style-type: none"> ▪ Console session timeout ▪ Logging level ▪ Audit implementation 	“How to Change Oracle Java Web Console Properties” on page 66

Getting Started With the Oracle Java Web Console

The Oracle Java Web Console's launch page lists the registered system management applications that you have permission to use, and displays a brief description of each application. You connect to a specific application by clicking its application name, which is a link to the actual application. By default, the selected application opens in the web console window. You can choose to open applications in separate browser windows by clicking the Start Each Application in a New Window check box. When you open applications in separate windows, the web console launch page remains available, so you can return to it and launch multiple applications under a single login.

To access the console launch page, type a URL of the following format in the web location field:

https://hostname.domain:6789

where the following applies:

- `https` specifies a Secure Socket Layer (SSL) connection
- `hostname.domain` specifies the name and domain of the server that is hosting the console
- `6789` is the console's assigned port number

Note – The first time you access the Oracle Java Web Console from a particular system, you must accept the server's certificate before the web console's launch page is displayed.

If RBAC is enabled on the system, and your user identity is assigned to a role, you are prompted for a role password after you have successfully logged in. If you assume a role, authorization checks are made for the assumed role. You can opt out of assuming a role by selecting NO ROLE, and then authorization checks are made against your user identity. Following a successful authorization check, the web console launch page is displayed.

▼ **How to Start Applications From the Oracle Java Web Console's Launch Page**

- 1 Start a web browser that is compatible with the Oracle Java Web Console, such as Mozilla 1.7 or Firefox 1.0.**

See “Supported Web Browsers” on page 58 for a list of supported browsers.

- 2 Type the console's URL in the web browser's location field.**

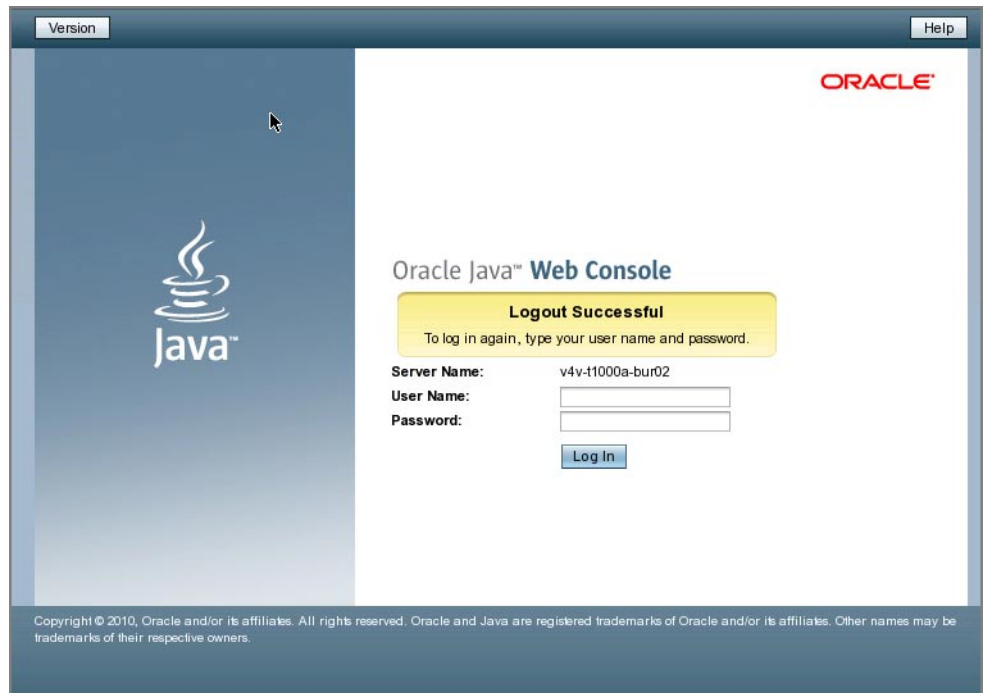
For example, if the management server host is named `sailfish`, and the domain is `sw`, the URL is `https://sailfish.sw:6789`. This URL takes you to the web console login page.

- 3 Accept the server's certificate.**

You only have to accept the server's certificate once per browser session, not each time you login to the console or start an application.

The login page is displayed as shown in the following figure.

FIGURE 3-1 Oracle Java Web Console Login Page



4 Enter your user name and password, and optionally your RBAC role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

The console services check your credentials to authenticate them, and ensure that you are authorized to use the console and registered applications.

5 If you want to run the application in a new window, click the Start Each Application in a New Window check box.

If you do not select this option, the application will run in the default window, replacing the launch page.

6 Click the link for the application that you want to run.

Tip – You can also launch an individual application directly and bypass the launch page by using the following syntax:

```
https://hostname.domain:6789/app-context-name
```

where *app-context-name* is the name that is used when the application is deployed.

To find the application context name, you can do one of the following:

- Read the application's documentation.
 - Run the `wcadmin list -a` or the `smreg list -a` command to see a list of deployed web applications and their context names.
 - Run the application from the web console's launch page and note the URL that is displayed in the address location field. You can type the URL directly the next time you use the application. Or, you can bookmark the location and access the application through the bookmark.
-

Managing the Console Service

Solaris 10 11/06: The Oracle Java Web Console service is managed through the Service Management Facility (SMF). You can start, stop, enable, and disable the console service by using SMF commands, or by using the `smcwebserver` script. The FMRI used in SMF for the console is `system/webconsole:console`.

▼ How to Start the Console Service

This procedure starts the server temporarily. If the server was disabled from starting when the system boots, it will continue to be disabled. If the server was enabled, it will continue to be enabled.

Starting with the Solaris 10 11/06 release, the running enabled state displays as `true` (temporary), if the server is running while disabled.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Start the server now, without changing the enabled state.

```
# smcwebserver start
```

▼ How to Enable the Console Service to Run at System Start

This procedure enables the console service to run each time the system starts. The console is not started in the current session.

Starting with the Solaris 10 11/06 release, this procedure sets the `general/enabled` property to `true` in SMF, so that the server is started at the time the system boots.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Enable the server to be started at system boot.

```
# smcwebserver enable
```

Solaris 10 11/06: Alternatively, if you want to both start the server now, and enable the server to start when the system boots, use the command:

```
# svcadm enable system/webconsole:console
```

Note – If you are running the Solaris 10 11/06 release, you cannot enable the console by using the `smcwebserver` command. You must use the `svcadm` command.

▼ How to Stop the Console Service

This procedure stops the server temporarily. If the server is disabled from starting when the system boots, it will continue to be disabled. If the server was enabled, it will continue to be enabled.

Starting with the Solaris 10 11/06 release, the running enabled state displays as `false` (temporary) if the server is stopped while enabled.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Stop the server now, without changing the enabled state.

```
# smcwebserver stop
```

▼ How to Disable the Console Service

When the console server is disabled, the server does not start when the system boots.

Starting with the Solaris 10 11/06 release, this procedure sets the console's `general/enabled` property to `false` in SMF, so that the console server does not start when the system boots.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Disable the server from starting when the system boots.

```
# smcwebserver disable
```

Solaris 10 11/06: Alternatively, if you want to both stop the server now, and disable the server from starting when the system boots, use the command:

```
# svcadm disable system/webconsole:console
```

Note – If you are running the Solaris 10 11/06 release, you cannot disable the console with the `smcwebserver` command. You must use the `svcadm` command.

Configuring the Oracle Java Web Console

The Oracle Java Web Console is preconfigured to run without administrator intervention. However, you might choose to change some of the web console's default behavior by overriding the console's configuration properties.

Note – Starting with the Solaris 10 11/06 OS, you must use the `wcadmin` command to change these properties. Previously, the `smreg` command was used. For more information about the `wcadmin` command, see the `wcadmin(1M)` man page.

Properties in the console's configuration files control the behavior of the console. To change the behavior, you define new values for properties to override the default values. The default values of most properties should not be overridden unless there is a specific need that the default values do not provide, such as specifying your own login service.

In general, the property values that you might consider changing are the following:

- **Console session timeout**

The web console's session timeout period is controlled by the `session.timeout.value` property. This property controls how long a web console page can be displayed without user interaction before the session times out. After the timeout is reached, the user must log in again. The default value is 15 minutes. You can set a new value, in minutes, to conform to your own security policy. However, keep in mind that this property controls the timeout period for all console users and all registered applications.

See [Example 3–1](#) for an example of how to change the session timeout.

- **Logging level**

You use logging properties to configure the logging service. The console log files are created in the `/var/log/webconsole/console` directory. The `logging.default.level` property determines which messages are logged. The console logs provide valuable information for troubleshooting problems.

The logging level applies to any messages that are written through the logging service, which by default uses syslog in the Oracle Solaris release. The syslog log file is `/var/adm/messages`. The file `/var/log/webconsole/console/console_debug_log` contains log messages written when the debugging service is enabled. This is done by setting the `debug.trace.level` property as described in “Using the Console Debug Trace Log” on [page 69](#). Although the default logging and debug logging services are separate, all Oracle Java Web Console logging messages to syslog are also written to the `console_debug_log` to aid in debugging. Generally, the logging service, set with `logging.default.level`, should be always enabled for logging by console applications. Debug logging, set with `debug.trace.level`, should only be enabled to investigate problems.

The following property values are available for `logging.default.level`:

- all
- info
- off
- severe
- warning

See [Example 3–2](#) for an example that shows how to change the logging level.

▪ **Auditing implementation**

Auditing is the process of generating and logging security-related management events. An event signifies that a specific user has updated the management information on a system. The auditing implementation is used by services and applications that generate audit events.

The following audit events are defined by the web console:

- Login
- Logout
- Role assumption

When audit events occur, a record of the event is made in an audit log. The location of the audit log varies with the auditing implementation that is in use. The web console's auditing service uses an auditing implementation that is provided by the underlying operating system.

The web console supports three auditing implementations: `Solaris`, `Log`, and `None`. You can select an auditing implementation by specifying one of these keywords for the value of the `audit.default.type` configuration property. Only one auditing implementation is in effect at a time.

The supported auditing implementation types are:

- `Solaris`

The `Solaris` implementation is the default. This implementation supports the BSM auditing mechanism. The auditing mechanism writes audit records into a system file in the `/var/audit` directory.

You can display the records with the `praudit` command. For events to be captured, you must enable the BSM auditing mechanism on the system. In addition, the `/etc/security/audit_control` file must contain entries that indicate which events should be generated. You must set the `lo` event as the `flag` option to see login and logout events for each user. For more information, see the `praudit(1M)` and `bsmconv(1M)` man pages and [Part VII, “Oracle Solaris Auditing,”](#) in *System Administration Guide: Security Services*.

- **Log**
You can configure this implementation to write to the system's `syslog` service. Audit messages are written to the console log if the logging service has been enabled at the `info` level. See [Example 3–2](#) for more information.
- **None**
No audit events are generated. Audit messages are written to the debug trace log, if enabled.

See [Example 3–5](#) for an example of specifying the auditing implementation.

▼ How to Change Oracle Java Web Console Properties

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.

2 Depending on which Oracle Solaris release you are running, change the selected property value as follows:

- **If you are running at least the Solaris 10 11/06 release, use this command:**

```
# wadmin add -p -a console name=value
```

`-p` Specifies that the object type is a property.

`-a console` Specifies that the property changes are for the application named `console`. The `-a console` option must always be used when you are changing console properties.

`name=value` Specifies the property name and the new value for that property.

- **For the Oracle Solaris 10, Solaris 10 1/06, and the Solaris 10 6/06 releases, use this command:**

```
# smreg add -p -c name
```

3 (Optional) Reset a console property to its default value.

- If you are running at least the Solaris 10 11/06 release, use this command:

```
# wcadmin remove -p -a console name=value
```

- For the Oracle Solaris 10, Solaris 10 1/06, and the Solaris 10 6/06 releases, use this command:

```
# smreg remove -p -c name
```

-p Specifies that the object type is a property.

-c Specifies that the property changes are for the console application. The -c option must always be used when you are changing console properties.

name Specifies the property name and the new value for that property.

Example 3-1 Changing the Oracle Java Web Console's Session Timeout Property

This example shows how to set the session time out value to 5 minutes.

```
# wcadmin add -p -a console session.timeout.value=5
```

Example 3-2 Configuring the Oracle Java Web Console Logging Level

This example shows you how to set the logging level to all.

```
# wcadmin add -p -a console logging.default.level=all
```

Example 3-3 Resetting the Oracle Java Web Console Logging Level to the Default Value

This example shows how to reset the logging level to the default.

```
# wcadmin remove -p -a console logging.default.level
```

Example 3-4 Specifying a Java Version for the Oracle Java Web Console

This example shows how to set the Java version for the console.

```
# wcadmin add -p -a console java.home=/usr/java
```

Example 3-5 Choosing an Auditing Implementation for the Oracle Java Web Console

This example shows you how to set the auditing implementation to None.

```
# wcadmin add -p -a console audit.default.type=None
```

The valid auditing types are:

None	No auditing
Log	Audit messages to syslog
Solaris	Audit messages to BSM

Oracle Java Web Console User Identity

By default, the web console runs under the user identity, `noaccess`. However, some system configurations disable the `noaccess` user, or set the login shell for the `noaccess` user to an invalid entry to make this user identity unusable.

When the `noaccess` user is not usable, the web console server cannot be started or configured, so an alternative user identity must be specified. Ideally, the user identity should be changed only once, before the console server is configured at initial startup.

You can configure the web console to run under an alternative non-root user identity by using either of the following commands before the console starts:

```
# smcwebserver start -u username
```

This command starts the web console server under the specified user identity. The web console server runs under this identity each time the server is subsequently started if the command is issued before the first console start.

If you are running at least the Solaris 10 11/06 release, you can also use this command:

```
# wadmin add -p -a console com.sun.web.console.user=  
username
```

Note – Starting with the Solaris 10 11/06 release, when the system initially starts, the console also starts and is automatically configured to run under `noaccess`. Consequently, the user identity is set to `noaccess` before you are able to change the user identity. Use the following commands to reset the console to its initial unconfigured state. Then, specify a different user identity when you restart the console.

```
# smcwebserver stop  
# /usr/share/webconsole/private/bin/wcremove -i console  
# smcwebserver start -u new_user_identity
```

For the Oracle Solaris 10, Solaris 10 1/06, Solaris 10 6/06 releases, use this command:

```
# smreg add -p -c com.sun.web.console.user=username
```

This command causes the web console server to run under the specified user identity the next time the server starts, and each time the server is started.

Using the Console Debug Trace Log

By default, the console does not log debug messages. You can turn on debug logging to help troubleshoot console service problems.

Use the `debug.trace.level` property to turn on debug logging by setting the property to a value other than 0.

Available choices include the following:

- 1 - Use this setting to record potentially severe errors.
- 2 - Use this setting to record important messages, as well as error messages of the 1 level.
- 3 - Use this setting to record all possible messages with full details.

By default, the debug trace log is created in the `/var/log/webconsole` directory for the Oracle Solaris 10, Solaris 10 1/06, and the Solaris 10 6/06 releases. **Starting with the Solaris 10 11/06 release**, the log is created in the `/var/log/webconsole/console` directory. The log file is named `console_debug_log`. Historical logs, such as `console_debug_log.1` and `console_debug_log.2` might also exist in this directory. There can be up to five (default setting) historical logs stored in this directory before the earliest log is deleted and a new log is created.

EXAMPLE 3-6 Setting the Console Debug Trace Log Level

Use the following command to set the debug trace log level to 3.

For the **Solaris 10 11/06** release, use this command:

```
# wadmin add -p -a console.debug.trace.level=3
```

For the Oracle Solaris 10, Solaris 10 1/06, and the Solaris 10 6/06 releases, use this command:

```
# smreg add -p -c debug.trace.level=3
```

EXAMPLE 3-7 Checking the Status of the debug.trace.level Property

To check the status of the `debug.trace.level` property, use the `wadmin list` or `smreg list` command.

Solaris 10 11/06:

EXAMPLE 3-7 Checking the Status of the `debug.trace.level` Property *(Continued)*

```
# wadmin list -p | grep "debug.trace.level"
```

For the Oracle Solaris 10, Solaris 10 1/06, and Solaris 10 6/06 releases, use this command:

```
# smreg list -p | grep "debug.trace.level"
```

Troubleshooting the Oracle Java Web Console Software (Task Map)

Task	Description	For Instructions
Check to determine if the console is running and enabled.	Use the <code>smcwebserver</code> , <code>wadmin</code> , and <code>svcs</code> commands to check if the console is running and enabled. This information is useful for troubleshooting problems.	“How to Check if the Console is Running and Enabled” on page 71
List console resources and properties.	You might need to gather information about the console resources and properties for troubleshooting purposes.	“How to List Console Resources and Properties” on page 72
Determine if an application is a legacy application.	Current applications are registered and deployed with a single command while the console server is running. Legacy applications require the console server to be stopped during registration. If you need to register or unregister an application, you must first determine if the application is a legacy application.	“How to Determine if an Application is a Legacy Application” on page 74
List all registered applications.	You can list all applications that are registered with the Oracle Java Web Console. Listing all registered applications provides you with information that can be helpful in troubleshooting situations.	“How to List Deployed Applications” on page 74
Register a legacy application with the Oracle Java Web Console.	If you need to use a legacy application, you must first register the application with the Oracle Java Web Console.	“How to Register a Legacy Application With the Oracle Java Web Console” on page 75

Task	Description	For Instructions
Unregister a legacy application from the Oracle Java Web Console.	If you do not want a legacy application registered with the Oracle Java Web Console, follow the procedure to unregister the legacy application.	“How to Unregister a Legacy Application From the Oracle Java Web Console” on page 76
Register a current application with the Oracle Java Web Console.	Before using a new application, you need to register the application with the Oracle Java Web Console.	“How to Register a Current Application With the Oracle Java Web Console” on page 77
Unregister a current application from the Oracle Java Web Console.	In some situations, you might need to unregister a current application from the Oracle Java Web Console.	“How to Unregister a Current Application from the Oracle Java Web Console” on page 78
Enable remote Access to the Oracle Java Web Console.	You can enable remote access only to the console, while leaving the other access restrictions in place.	“How to Enable Remote Access to the Oracle Java Web Console” on page 82
Change the console's internal passwords.	The Oracle Java Web Console uses internal passwords. To reduce the possibility of a security breach, you can change these passwords.	“How to Change the Console's Internal Passwords” on page 83

Troubleshooting the Oracle Java Web Console Software

The following information is provided to help you troubleshoot any problems that you might encounter when using the Oracle Java Web Console software.

Checking Console Status and Properties

You can use the `smcwebserver`, `wcadmin`, and `svcs` commands to get different types of information about the console, which might be useful for troubleshooting problems.

▼ How to Check if the Console is Running and Enabled

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Check the server status.

```
# smcwebserver status
Sun Java(TM) Web Console is running
```

3 Solaris 10 11/06: Check the console's SMF status and enabled state.

```
# svcs -l system/webconsole:console
```

If you start and stop the server with `smcwebserver` commands without enabling and disabling, the enabled property might display as `false` (temporary) or `true` (temporary).

▼ How to List Console Resources and Properties

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.

2 List the console's resources and properties.

If you are running at least the Solaris 10 11/06 release, use this command:

```
# wadmin list
```

Deployed web applications (application name, context name, status):

```
console      ROOT          [running]
console      com_sun_web_ui [running]
console      console       [running]
console      manager       [running]
legacy      myapp         [running]
```

Registered jar files (application name, identifier, path):

```
console audit_jar      /usr/lib/audit/Audit.jar
console console_jars /usr/share/webconsole/lib/*.jar
console jato_jar       /usr/share/lib/jato/jato.jar
console javahelp_jar /usr/jdk/packages/javahelp-2.0/lib/*.jar
console shared_jars /usr/share/webconsole/private/container/shared/lib/*.jar
```

Registered login modules (application name, service name, identifier):

```
console ConsoleLogin userlogin
console ConsoleLogin rolelogin
```

Shared service properties (name, value):

```
ENABLE      yes
java.home   /usr/jdk/jdk1.5.0_06
```

Note – This ENABLE property is ignored because SMF uses its own enabled property, which is shown in the previous procedure. The ENABLE property is used on older systems where the console server is not managed by SMF.

For the Oracle Solaris 10, Solaris 10 1/06, and Solaris 10 6/06 releases, use this command:

```
# smreg list
```


The list of registered plugin applications:

```
com.sun.web.console_2.2.4      /usr/share/webconsole/console
com.sun.web.ui_2.2.4         /usr/share/webconsole/com_sun_web_ui
com.sun.web.admin.example_2.2.4 /usr/share/webconsole/example
```

The list of registered jar files:

```
com_sun_management_services_api.jar scoped to ALL
com_sun_management_services_impl.jar scoped to ALL
com_sun_management_console_impl.jar scoped to ALL
com_sun_management_cc.jar scoped to ALL
com_sun_management_webcommon.jar scoped to ALL
com_ipplanet_jato_jato.jar scoped to ALL
com_sun_management_solaris_impl.jar scoped to ALL
com_sun_management_solaris_implx.jar scoped to ALL
```

The list of registered login modules for service ConsoleLogin:

```
com.sun.management.services.authentication.PamLoginModule optional
use_first_pass="true" commandPath="/usr/lib/webconsole";
com.sun.management.services.authentication.RbacRoleLoginModule requisite
force_role_check="true" commandPath="/usr/lib/webconsole";
```

The list of registered server configuration properties:

```
session.timeout.value=15
authentication.login.cliservice=ConsoleLogin
logging.default.handler=com.sun.management.services.logging.ConsoleSyslogHandler
logging.default.level=info
logging.default.resource=com.sun.management.services.logging.resources.Resources
logging.default.filter=none
logging.debug.level=off
audit.default.type=None
audit.None.class=com.sun.management.services.audit.LogAuditSession
audit.Log.class=com.sun.management.services.audit.LogAuditSession audit.class.fail=none
authorization.default.type=SolarisRbac
authorization.SolarisRbac.class=
com.sun.management.services.authorization.SolarisRbacAuthorizationService
authorization.PrincipalType.class=
com.sun.management.services.authorization.PrincipalTypeAuthorizationService
debug.trace.level=0
.
.
.
No environment properties have been registered.
```

Problems Accessing the Console

Problems with console access might indicate that the console server is not enabled, or security settings are restrictive. See [“Checking Console Status and Properties” on page 71](#) and [“Oracle Java Web Console Security Considerations” on page 78](#) for more information.

Problems with Application Registration

This section contains information about solving possible registration problems with console applications. For information about a particular console application, you should refer to the application's documentation.

Note – Console applications typically are registered as part of their installation process, so you should not normally need to register an application yourself.

Starting with the Solaris 10 11/06 release, the web console has changed the approach to application registration but can still support applications that were developed for earlier versions of the console. Current applications are registered and deployed with a single command while the console server is running. Applications that were developed for the earlier console are known as *legacy* applications, and require the console server to be stopped during registration. If you need to register or unregister an application, you must first determine if the application is a legacy application, as described in the following procedure.

▼ How to Determine if an Application is a Legacy Application

1 View the application's `app.xml` file.

The `app.xml` file is located in the application's `WEB-INF` directory.

2 Examine the `registrationInfo` tag in the `app.xml` file.

For a legacy application, the `registrationInfo` tag is a version 2.x. For example, `registrationInfo version="2.2.4"`.

For a current application, the version in the `registrationInfo` tag is at least 3.0. For example, `registrationInfo version="3.0"`.

▼ How to List Deployed Applications

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

2 List the deployed applications.

If you are running at least the Solaris 10 11/06 release, use this command:

```
# wadmin list -a
```

Deployed web applications (application name, context name, status):

```
console ROOT [running]
```

```

console  com_sun_web_ui  [running]
console  console          [running]
console  manager          [running]
legacy   myapp            [running]

```

The command lists all the registered and deployed applications. Legacy applications are listed with the application name `legacy`. See [“How to Determine if an Application is a Legacy Application” on page 74](#). All other listed applications are current applications, and would be registered as described in [“How to Register a Current Application With the Oracle Java Web Console” on page 77](#).

Typically, the status that is shown for the applications contains either `running` or `stopped`. If the status is `running`, the application is currently loaded and available. If the status is `stopped`, then the application is not currently loaded and is unavailable. Sometimes an application registers and deploys successfully, but does not load because of a problem in the application. If so, the application's status is `stopped`. Check the `console_debug_log` to determine if there is an error with a traceback from the console's underlying web container, Tomcat, when attempting to load the application. For more information about the `console_debug_log`, see [“Using the Console Debug Trace Log” on page 69](#).

If all the applications show `stopped` (including the console application), this usually means the console's web container is not running. The list of applications in this case is obtained from the static context.xml files registered with the web container.

For the Oracle Solaris 10, Solaris 10 1/06, and Solaris 10 6/06 releases, use this command:

```
# smreg list -a
```

The list of registered plugin applications:

```

com.sun.web.console_2.2.4      /usr/share/webconsole/console
com.sun.web.ui_2.2.4          /usr/share/webconsole/com_sun_web_ui
com.sun.web.admin.yourapp_2.2.4 /usr/share/webconsole/yourapp

```

▼ How to Register a Legacy Application With the Oracle Java Web Console

Note – This procedure applies to all console applications in the Oracle Solaris 10, Solaris 10 1/06, and Solaris 10 6/06 releases. **Starting with Solaris 10 11/06 release**, this procedure also applies *only* to those applications that are identified as legacy applications. See [“How to Register a Current Application With the Oracle Java Web Console” on page 77](#) for the registration procedure for current applications. See also [“How to Determine if an Application is a Legacy Application” on page 74](#).

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Stop the web server.

```
# smcwebserver stop
```

3 Register the application.

```
# smreg add -a /directory/containing/application-files
```

The smreg command manages the information in the Oracle Java Web Console's registration table. This script also performs some additional work to deploy the application. For additional options to this command, see the [smreg\(1M\)](#) man page.

4 Restart the web server.

```
# smcwebserver start
```

Example 3–8 Registering a Legacy Application

This example shows how to register a legacy application whose files are located in the `/usr/share/webconsole/example` directory. Notice that for legacy applications, the console server must be stopped before the application is registered, and started after the application is registered. A warning given by smreg can be ignored because this application is a legacy console application.

```
# smcwebserver stop
# smreg add -a /usr/share/webconsole/example
```

```
Warning: smreg is obsolete and is preserved only for
compatibility with legacy console applications. Use wadmin instead.
```

```
Type "man wadmin" or "wadmin --help" for more information.
```

```
Registering com.sun.web.admin.example_version.
```

```
# smcwebserver start
```

▼ **How to Unregister a Legacy Application From the Oracle Java Web Console**

Note – This procedure applies to all console applications in the Oracle Solaris 10, Solaris 10 1/06, and Solaris 10 6/06 releases. **Starting with Solaris 10 11/06 release**, this procedure applies *only* to those applications that are identified as legacy applications. See [“How to Unregister a Current Application from the Oracle Java Web Console” on page 78](#) for the procedure that describes how to unregister current applications.

If you do not want a particular legacy application to display in the web console's launch page, but you do not want to uninstall the software, you can use the smreg command to unregister the application. See [“How to Determine if an Application is a Legacy Application” on page 74](#).

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Unregister an application.

```
# smreg remove -a app-name
```

Example 3–9 Unregistering a Legacy Application From the Oracle Java Web Console

This example shows how to unregister a legacy application with the *app-name* `com.sun.web.admin.example`.

```
# smreg remove -a com.sun.web.admin.example

Unregistering com.sun.web.admin.example_version.
```

▼ **How to Register a Current Application With the Oracle Java Web Console**

Solaris 10 11/06: This procedure is for updated console applications that can be registered and deployed without stopping and starting the console server. See “How to Register a Legacy Application With the Oracle Java Web Console” on page 75 for the registration procedure for legacy applications and all console applications that are in the Oracle Solaris 10, Solaris 10 1/06, Solaris 10 6/06 releases. See also “How to Determine if an Application is a Legacy Application” on page 74.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Register and deploy the application.

```
wcadmin deploy -a app-name -x app-context-name
/full path/to/app-name
```

Example 3–10 Registering Current Applications

This example shows how to register and deploy an application that has been developed or updated for the current web console.

```
# wcadmin deploy -a newexample_1.0 -x newexample /apps/webconsole/newexample
```

▼ How to Unregister a Current Application from the Oracle Java Web Console

Solaris 10 11/06: This procedure is for updated console applications, which can be unregistered and undeployed without stopping and starting the console server. See [“How to Unregister a Legacy Application From the Oracle Java Web Console” on page 76](#) for the unregistration procedure for legacy applications and all console applications that are in the Oracle Solaris 10, Solaris 10 1/06, Solaris 10 6/06 releases. See [“How to List Deployed Applications” on page 74](#) and [“How to Determine if an Application is a Legacy Application” on page 74](#) to determine if an application is a legacy or updated application.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Undeploy and unregister the application.

```
# wadmin undeploy -a newexample_1.0 -x newexample
```

Oracle Java Web Console Reference Information

This reference section includes the following topics:

- [“Oracle Java Web Console Security Considerations” on page 78](#)
- [“Specifying Authorizations With the authTypes Tag” on page 80](#)

Oracle Java Web Console Security Considerations

There are several security considerations to keep in mind when you use applications that are in the Oracle Java Web Console.

These security considerations include the following:

- **Access to the Oracle Java Web Console** – Whether you can connect to the console through a browser.
- **Access to applications** – Whether you can see a particular application in the Oracle Java Web Console's launch page.
- **Application permissions** – The levels of permissions that you must have to run parts or all of an application.
- **Application access to remote systems** – How security credentials relate to remote systems.
- **Internal passwords used in the console** – Changing the default passwords that are used internally in the console, starting with the Solaris 10 11/06 release.

Access to the Oracle Java Web Console

Permissions to the web console launcher application are usually open so that any valid user can log in. However, you can restrict access to the console by specifying the rights in the `authTypes` tag in the web console's `app.xml` file, which is located in the `/usr/share/webconsole/webapps/console/WEB-INF` directory. For more information, see [“Specifying Authorizations With the `authTypes` Tag” on page 80](#).

Some system configurations are set up to be very secure, so that attempts to connect from a remote system to the URLs of the console or registered applications are refused. If your system is configured to prevent remote access, when you try to access the console as `https://hostname.domain:6789`, your browser displays a message such as:

```
Connect to hostname.domain:6789 failed (Connection refused)
```

The SMF profile in effect on the system might be restricting access. See [“SMF Profiles” on page 352](#) for more information about profiles. See [“Enabling Remote Access to the Oracle Java Web Console” on page 82](#) for a procedure to allow access to the console from remote systems.

Access to Applications in the Oracle Java Web Console

After you successfully log in to the web console, you might not automatically have access to all of the applications that are registered in that console. Typically, applications are installed so that all users can see them in the console launch page. As an administrator, you can grant and restrict access to applications.

To restrict access to an application, specify the rights in the `authTypes` tag, which is in the application's `app.xml` file. You can find the application's `app.xml` file in the `installation-location/WEB-INF/` subdirectory. Typically, this directory would be located in `/usr/share/webconsole/webapps/ app-context-name/WEB-INF`.

If the application files are not in the usual location, you can locate the files by using the following command:

```
wcadmin list --detail -a
```

This command lists each deployed application, showing when it was deployed and the path to the application's base directory. The `app.xml` file is located in the subdirectory `WEB-INF` within the base directory.

For more information, see [“Specifying Authorizations With the `authTypes` Tag” on page 80](#).

Application Privileges

If you can see an application's link on the Oracle Java Web Console's launch page, you can run that application. However, an application might make additional authorization checks based upon the authenticated user or role identity. These checks are not controlled by the `authTypes` tag, but are explicitly coded into the application itself. For example, an application might grant read access to all authenticated users, but restrict update access to a few users or a few roles.

Application Access to Remote Systems

Having all the appropriate credentials does not guarantee that you can use an application to manage every system within the application's scope of operation. Each system that you administer by using the Oracle Java Web Console application has its own security domain. Having read-and-write permissions on the web console system does not guarantee that those credentials are automatically sufficient to administer any other remote system.

In general, access to remote systems depends on how the security is implemented in the web application. Typically, web applications make calls to *agents* that perform actions on behalf of the applications. These applications must be authenticated by the agents based on their web console credentials and the credentials by which they are known on the agent system. Depending upon how this agent authentication is done, an authorization check might also be made on the agent itself, based upon this authenticated identity.

For example, in web applications that use remote WBEM agents, authentication typically uses the user or role identity that initially authenticated to the Oracle Java Web Console. If this authentication fails on that agent system, access to that system is denied in the web application. If authentication succeeds on that agent system, access might still be denied if the agent makes an access control check and denies access there. Most applications are written so that the authentication and authorization checks on the agent never fail if you have been successfully authenticated on the web console and assumed the correct role.

Internal Passwords That Are Used in the Console

Starting with the Solaris 10 11/06 release, the Oracle Java Web Console uses several password-protected internal user names to perform administrative tasks on the underlying web server, and to encrypt key store and trust store files. The passwords are set to initial values to enable the console to be installed. To reduce the possibility of a security breach, you should change the passwords after installation. See [“Changing Internal Passwords for Oracle Java Web Console”](#) on page 83

Specifying Authorizations With the `authTypes` Tag

While most system management web applications do not require any administrator intervention to use the `authTypes` tag, in some cases, you might need to change the values of this tag. The `authTypes` tag contains a set of information that describes the level of authorization that is required for a user to view an application in the Oracle Java Web Console. The web console determines if a user is authorized to see a particular application, based on the authorization requirements in the application's `app.xml` file. Each application can determine whether a user must have proper authorization to run the application. This determination might be made as part of the application installation process. Or, you might need to supply the information, depending on your own security requirements. The product documentation for the application should contain the information that is necessary to determine whether you need to specify a particular permission.

You can nest several `authType` tags within the `authTypes` tag.

The `authTypes` tag must contain at least one `authType` tag that provides the following necessary information:

- Type of authorization check to perform
- Permission subclass name
- Parameters that are required to instantiate the `Permission` subclass

In the following example, the `authType` tag has one attribute, `name`. The required `name` attribute is the name of the authorization service type. Different authorization types might require different values for the `classType` and `permissionParam` tags.

```
<authTypes>
  <authType name="SolarisRbac">
    <classType>
      com.sun.management.solaris.RbacPermission
    </classType>
    <permissionParam name="permission">
      solaris.admin.serialmgr.read
    </permissionParam>
  </authType>
</authTypes>
```

The following table shows the tags that can be nested within an `authType` tag.

TABLE 3-1 Nested `authType` Tags

Tag	Attribute	Description
<code>classType</code>		The <code>Permission</code> subclass name. This tag is required.
<code>permissionParam</code>	<code>name</code>	The parameters that are required to create an instance of the class specified by <code>classType</code> .

The `authTypes` tag and nested `authType` tags are required elements in the `app.xml` file. If you want to register an application that is available to anyone, specify the `authType` tag with no content, as shown in the following example.

```
<authTypes>
  <authType name="">
    <classType></classType>
    <permissionParam name=""></permissionParam>
  </authType>
</authTypes>
```

Enabling Remote Access to the Oracle Java Web Console

If you can only connect to the console by logging into the system that is running the console, and then using the URL, `https://localhost:6789`, the system is using a configuration that prevents remote access. **Starting with the Solaris 10 11/06 release**, you can enable remote access only to the console, while leaving the other access restrictions in place, by using the following procedure:

▼ How to Enable Remote Access to the Oracle Java Web Console

- 1 **Become superuser or assume an equivalent role on the system where the console is running.**

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

- 2 **Set a property to allow the console server to respond to network requests and restart the console server.**

```
# svccfg -s svc:/system/webconsole setprop options/tcp_listen = true

# smcwebservice restart
```

Disabling Remote Access to the Oracle Java Web Console

You can prevent users from connecting to the console from remote systems. **Starting with the Solaris 10 11/06 release**, you can disable remote access only to the console, while leaving the other access permissions in place, by using the following procedure:

▼ How to Disable Remote Access to the Oracle Java Web Console

- 1 **Become superuser or assume an equivalent role on the system where the console is running.**

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

- 2 **Set a property to prevent the console server from responding to network requests, and restart the console server.**

```
# svccfg -s svc:/system/webconsole setprop options/tcp_listen = false

# smcwebservice restart
```

After the restart the console now only responds to a browser on the same system as the console server process. You cannot use a proxy in the browser, only a direct connection. You can also use the `https://localhost:6789/` URL to access the console.

Changing Internal Passwords for Oracle Java Web Console

Starting with the Solaris 10 11/06 release, the console uses some internal user names and passwords. The console's internal user names and passwords are used only by the console framework, and are never used directly by a user or system administrator. However, if the passwords were known, a malicious user could potentially interfere with the console applications. To reduce the possibility of such a security breach, you should change the passwords. You do not need to remember the new passwords, because the software uses them invisibly.

▼ How to Change the Console's Internal Passwords

The passwords are known as the administrative password, keystore password, and truststore password. You do not need to know the default initial values in order to change the passwords. This procedure explains how to change all three passwords with separate commands.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Change the administrative password.

```
# wadmin password -a
```

You are prompted to enter the new password twice. The password should be 8 to 32 characters.

3 Change the key store password.

```
# wadmin password -k
```

You are prompted to enter the new password twice. The password should be 8 to 32 characters.

4 Change the trust store password.

```
# wadmin password -t
```

You are prompted to enter the new password twice. The password should be 8 to 32 characters.

Managing User Accounts and Groups (Overview)

This chapter provides guidelines and planning information for managing user accounts and groups. This chapter also includes information about customizing the user's work environment.

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

- “What's New or Changed in Managing Users and Groups?” on page 85
- “What Are User Accounts and Groups?” on page 86
- “Where User Account and Group Information Is Stored” on page 93
- “Tools for User Account and Group Account Management” on page 86
- “Customizing a User's Work Environment” on page 101

For step-by-step instructions on managing user accounts and groups, see [Chapter 5, “Managing User Accounts and Groups \(Tasks\)”](#).

What's New or Changed in Managing Users and Groups?

This section includes information about new or changed features for managing users and groups in this Oracle Solaris release.

In this Solaris release, there are no new or changed features.

For a complete listing of new features and a description of Oracle Solaris releases, see [Oracle Solaris 10 8/11 What's New](#).

Tools for User Account and Group Account Management

The following table describes available tools for user account and group management.

TABLE 4-1 Tools for User Account and Group Management

Tool Name	Description	For More Information
Solaris Management Console	Graphical tool that is used to manage users, groups, roles, rights, mailing lists, disks, terminals, and modems.	“Setting Up User Accounts (Task Map)” on page 113
smuser, smrole, smgroup	Commands that are used to manage users, groups and roles. The SMC services must be running to use these commands.	“Adding a Group and User With the smgroup and smuser Commands” on page 119
useradd, groupadd, roleadd; usermod, groupmod, rolemod; userdel, groupdel, roledel	Commands that are used to manage users, groups, and roles.	“Adding a Group and User With the groupadd and useradd Commands” on page 119

You can add a group with the following tools:

- Solaris Management Console's Groups tool
- Admintool

Note – The Admintool is not available in this Solaris release.

TABLE 4-2 Solaris User/Group Command Descriptions

Command	Description	References
useradd, usermod, userdel	Adds, modifies, or removes a user	useradd(1M) , usermod(1M) , userdel(1M)
groupadd, groupmod, groupdel	Adds, modifies, or removes a group	groupadd(1M) , groupmod(1M) , groupdel(1M)

What Are User Accounts and Groups?

One basic system administration task is to set up a user account for each user at a site. A typical user account includes the information a user needs to log in and use a system, without having the system's root password. The components of user account information are described in [“User Account Components” on page 87](#).

When you set up a user account, you can add the user to predefined groups of users. A typical use of groups is to set up group permissions on a file and directory, which allows access only to users who are part of that group.

For example, you might have a directory containing confidential files that only a few users should be able to access. You could set up a group called `topsecret` that includes the users working on the `topsecret` project. And, you could set up the `topsecret` files with read permission for the `topsecret` group. That way, only the users in the `topsecret` group would be able to read the files.

A special type of user account, called a *role*, is used to give selected users special privileges. For more information, see [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services*.

User Account Components

The following sections describe the specific components of a user account.

User (Login) Names

User names, also called *login names*, let users access their own systems and remote systems that have the appropriate access privileges. You must choose a user name for each user account that you create.

Consider establishing a standard way of assigning user names so that they are easier for you to track. Also, names should be easy for users to remember. A simple scheme when selecting a user name is to use the first name initial and first seven letters of the user's last name. For example, Ziggy Ignatz becomes `zignatz`. If this scheme results in duplicate names, you can use the first initial, middle initial, and the first six characters of the user's last name. For example, Ziggy Top Ignatz becomes `ztignatz`.

If this scheme still results in duplicate names, consider using the following scheme to create a user name:

- The first initial, middle initial, first five characters of the user's last name
- The number 1, or 2, or 3, and so on, until you have a unique name

Note – Each new user name must be distinct from any mail aliases that are known to the system or to a NIS domain. Otherwise, mail might be delivered to the alias rather than to the actual user.

For detailed guidelines on setting up user (login) names, see [“Guidelines for Using User Names, User IDs, and Group IDs”](#) on page 93.

User ID Numbers

Associated with each user name is a user identification number (UID). The UID number identifies the user name to any system on which the user attempts to log in. And, the UID number is used by systems to identify the owners of files and directories. If you create user accounts for a single individual on a number of different systems, always use the same user name and ID number. In that way, the user can easily move files between systems without ownership problems.

UID numbers must be a whole number that is less than or equal to 2147483647. UID numbers are required for both regular user accounts and special system accounts. The following table lists the UID numbers that are reserved for user accounts and system accounts.

TABLE 4-3 Reserved UID Numbers

UID Numbers	User or Login Accounts	Description
0 – 99	root, daemon, bin, sys, and so on	Reserved for use by OracleSolaris
100 – 2147483647	Regular users	General purpose accounts
60001 and 65534	nobody and nobody4	Anonymous users
60002	noaccess	Non trusted users

Do not assign UIDs 0 through 99. These UIDs are reserved for allocation by Oracle Solaris. By definition, root always has UID 0, daemon has UID 1, and pseudo-user bin has UID 2. In addition, you should give uucp logins and pseudo user logins, such as who, tty, and ttytype, low UIDs so that they fall at the beginning of the `passwd` file.

For additional guidelines on setting up UIDs, see [“Guidelines for Using User Names, User IDs, and Group IDs” on page 93](#).

As with user (login) names, you should adopt a scheme for assigning unique UID numbers. Some companies assign unique employee numbers. Then, administrators add a number to the employee number to create a unique UID number for each employee.

To minimize security risks, you should avoid reusing the UIDs from deleted accounts. If you must reuse a UID, “wipe the slate clean” so that the new user is not affected by attributes set for a former user. For example, a former user might have been denied access to a printer by being included in a printer deny list. However, that attribute might be inappropriate for the new user.

Using Large User IDs and Group IDs

UIDs and group IDs (GIDs) can be assigned up to the maximum value of a signed integer, or 2147483647.

The following table describes interoperability issues with Oracle Solaris products and previous releases.

TABLE 4-4 Interoperability Issues for UIDs or GIDs Over 60000

Category	Product or Command	Issue
NFS interoperability	SunOS 4.0 NFS software and compatible releases	NFS server and client code truncates large UIDs and GIDs to 16 bits. This situation can create security problems, if systems running SunOS 4.0 and compatible releases are used in an environment where large UIDs and GIDs are being used. Systems running SunOS 4.0 and compatible releases require a patch to avoid this problem.
Name service interoperability	NIS name service and file-based name service	Users with UIDs greater than 60000 can log in or use the <code>su</code> command on systems running the Solaris 2.5 (and compatible releases). However, their UIDs and GIDs will be set to 60001 (<i>nobody</i>).
	NIS+ name service	Users with UIDs greater than 60000 are denied access on systems running Solaris 2.5 (and compatible releases) and the NIS+ name service.

The following table describes UID and GID limitations.

TABLE 4-5 Large UID and GID Limitation Summary

UID or GID	Limitations
262144 or greater	Users who use the <code>cpio</code> command with the default archive format to copy a file see an error message for each file. And, the UIDs and GIDs are set to <i>nobody</i> in the archive.
2097152 or greater	Users who use the <code>cpio</code> command with the <code>-H odc</code> format or the <code>pax -x cpio</code> command to copy files see an error message returned for each file. And, the UIDs and GIDs are set to <i>nobody</i> in the archive.
1000000 or greater	Users who use the <code>ar</code> command have their UIDs and GIDs set to <i>nobody</i> in the archive.
2097152 or greater	Users who use the <code>tar</code> command, the <code>cpio -H ustar</code> command, or the <code>pax -x tar</code> command have their UIDs and GIDs set to <i>nobody</i> .

UNIX Groups

A *group* is a collection of users who can share files and other system resources. For example, users who working on the same project could be formed into a group. A group is traditionally known as a UNIX group.

Each group must have a name, a group identification (GID) number, and a list of user names that belong to the group. A GID number identifies the group internally to the system.

The two types of groups that a user can belong to are as follows:

- **Primary group** – Specifies a group that the operating system assigns to files that are created by the user. Each user must belong to a primary group.
- **Secondary groups** – Specifies one or more groups to which a user also belongs. Users can belong to up to 15 secondary groups.

For detailed guidelines on setting up group names, see [“Guidelines for Using User Names, User IDs, and Group IDs” on page 93](#).

Sometimes, a user's secondary group is not important. For example, ownership of files reflect the primary group, not any secondary groups. Other applications, however, might rely on a user's secondary group memberships. For example, a user has to be a member of the `sysadmin` group (group 14) to use the `Admintool` software in previous Solaris releases. However, it doesn't matter if group 14 is his or her current primary group.

The `groups` command lists the groups that a user belongs to. A user can have only one primary group at a time. However, a user can temporarily change the user's primary group, with the `newgrp` command, to any other group in which the user is a member.

When adding a user account, you must assign a primary group for a user or accept the default group, `staff` (group 10). The primary group should already exist. If the primary group does not exist, specify the group by a GID number. User names are not added to primary groups. If user names were added to primary groups, the list might become too long. Before you can assign users to a new secondary group, you must create the group and assign it a GID number.

Groups can be local to a system or managed through a name service. To simplify group administration, you should use a name service such as NIS or a directory service such as LDAP. These services enable you to centrally manage group memberships.

User Passwords

You can specify a password for a user when you add the user. Or, you can force the user to specify a password when the user first logs in.

User passwords must comply with the following syntax:

- Password length must at least match the value identified by the `PASSLENGTH` variable in the `/etc/default/passwd` file. By default, `PASSLENGTH` is set to 6.
- The first 6 characters of the password must contain at least two alphabetic characters and have at least one numeric or special character.
- You can increase the maximum password length to more than eight characters by configuring the `/etc/policy.conf` file with an algorithm that supports greater than eight characters.

Although user names are publicly known, passwords must be kept secret and known only to users. Each user account should be assigned a password. The password can be a combination of six to eight letters, numbers, or special characters.

To make your computer systems more secure, users should change their passwords periodically. For a high level of security, you should require users to change their passwords every six weeks. Once every three months is adequate for lower levels of security. System administration logins (such as `root` and `sys`) should be changed monthly, or whenever a person who knows the root password leaves the company or is reassigned.

Many breaches of computer security involve guessing a legitimate user's password. You should make sure that users avoid using proper nouns, names, login names, and other passwords that a person might guess just by knowing something about the user.

Good choices for passwords include the following:

- Phrases (beammeup).
- Nonsense words made up of the first letters of every word in a phrase. For example, `swotr` for `SomeWhere Over The RainBow`.
- Words with numbers or symbols substituted for letters. For example, `sn00py` for `snoopy`.

Do not use these choices for passwords:

- Your name (spelled forwards, backwards, or jumbled)
- Names of family members or pets
- Car license numbers
- Telephone numbers
- Social Security numbers
- Employee numbers
- Words related to a hobby or interest
- Seasonal themes, such as `Santa` in December
- Any word in the dictionary

Home Directories

The home directory is the portion of a file system allocated to a user for storing private files. The amount of space you allocate for a home directory depends on the kinds of files the user creates, their size, and the number of files that are created.

A home directory can be located either on the user's local system or on a remote file server. In either case, by convention the home directory should be created as `/export/home/username`. For a large site, you should store home directories on a server. Use a separate file system for each `/export/homen` directory to facilitate backing up and restoring home directories. For example, `/export/home1`, `/export/home2`.

Regardless of where their home directory is located, users usually access their home directories through a mount point named `/home/username`. When AutoFS is used to mount home

directories, you are not permitted to create any directories under the /home mount point on any system. The system recognizes the special status of /home when AutoFS is active. For more information about automounting home directories, see [“Task Overview for Autofs Administration”](#) in *System Administration Guide: Network Services*.

To use the home directory anywhere on the network, you should always refer to the home directory as \$HOME, not as /export/home/*username*. The latter is machine-specific. In addition, any symbolic links created in a user's home directory should use relative paths (for example, ../../../../x/y/x) so that the links are valid no matter where the home directory is mounted.

Name Services

If you are managing user accounts for a large site, you might want to consider using a name or directory service such as LDAP, or NIS or NIS+. A name or directory service enables you to store user account information in a centralized manner instead of storing user account information in every system's /etc files. When you use a name or directory service for user accounts, users can move from system to system using the same user account without having site-wide user account information duplicated on every system. Using a name or directory service also promotes centralized and consistent user account information.

User's Work Environment

Besides having a home directory to create and store files, users need an environment that gives them access to the tools and resources they need to do their work. When a user logs in to a system, the user's work environment is determined by initialization files. These files are defined by the user's startup shell, which can vary, depending on the release.

A good strategy for managing the user's work environment is to provide customized user initialization files, such as .login, .cshrc, or .profile, in the user's home directory.

Note – Do not use system initialization files, such as /etc/profile or /etc/.login, to manage a user's work environment. These files reside locally on systems and are not centrally administered. For example, if AutoFS is used to mount the user's home directory from any system on the network, you would have to modify the system initialization files on each system to ensure a consistent environment whenever a user moved from system to system.

For detailed information about customizing user initialization files for users, see [“Customizing a User's Work Environment”](#) on page 101.

Another way to customize user accounts is through role-based access control (RBAC). See [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services* for more information.

Guidelines for Using User Names, User IDs, and Group IDs

User names, UIDs, and GIDs should be unique within your organization, which might span multiple domains.

Keep the following guidelines in mind when creating user or role names, UIDs, and GIDs:

- **User names** – They should contain from two to eight letters and numerals. The first character should be a letter. At least one character should be a lowercase letter.

Note – Even though user names can include a period (.), underscore (_), or hyphen (-), using these characters is not recommended because they can cause problems with some software products.

- **System accounts** – Do not use any of the user names, UIDs, or GIDs that are contained in the default `/etc/passwd` and `/etc/group` files. Do not use the UIDs and GIDs, 0-99. These numbers are reserved for allocation by the Oracle Solaris OS and should not be used by anyone. Note that this restriction also applies to numbers not currently in use.

For example, `gdm` is the reserved user name and group name for the GNOME Display Manager daemon and should not be used for another user. For a complete listing of the default `/etc/passwd` and `/etc/group` entries, see [Table 4-6](#) and [Table 4-7](#).

The `nobody` and `nobody4` accounts should never be used for running processes. These two accounts are reserved for use by NFS. Use of these accounts for running processes could lead to unexpected security risks. Processes that need to run as a nonroot user should use the `daemon` or `noaccess` accounts.

- **System account configuration** – The configuration of the default system accounts should never be changed. This includes changing the login shell of a system account that is currently locked. The only exception to this rule is the setting of a password and password aging parameters for the root account.

Where User Account and Group Information Is Stored

Depending on your site policy, user account and group information can be stored in your local system's `/etc` files or in a name or directory service as follows:

- The NIS+ name service information is stored in tables.
- The NIS name service information is stored in maps.
- The LDAP directory service information is stored in indexed database files.

Note – To avoid confusion, the location of the user account and group information is generically referred to as a *file* rather than as a *database*, *table*, or *map*.

Most user account information is stored in the `passwd` file. Password information is stored as follows:

- In the `passwd` file when you are using NIS or NIS+
- In the `/etc/shadow` file when you are using `/etc` files
- In the `people` container when you are using LDAP

Password aging is available when you are using NIS+ or LDAP, but not NIS.

Group information is stored in the `group` file for NIS, NIS+, and files. For LDAP, group information is stored in the `group` container.

Fields in the `passwd` File

The fields in the `passwd` file are separated by colons and contain the following information:

username:password:uid:gid:comment:home-directory:login-shell

For example:

```
kryten:x:101:100:Kryten Series 4000 Mechanoid:/export/home/kryten:/bin/csh
```

For a complete description of the fields in the `passwd` file, see the [passwd\(1\)](#) man page.

Default `passwd` File

The default `passwd` file contains entries for standard daemons. Daemons are processes that are usually started at boot time to perform some system-wide task, such as printing, network administration, or port monitoring.

```
root:x:0:0:Super-User:/:/sbin/sh
daemon:x:1:1:/:
bin:x:2:2:/:usr/bin:
sys:x:3:3:/:
adm:x:4:4:Admin:/var/adm:
lp:x:71:8:Line Printer Admin:/usr/spool/lp:
uucp:x:5:5:uucp Admin:/usr/lib/uucp:
nuucp:x:9:9:uucp Admin:/var/spool/uucppublic:/usr/lib/uucp/uucico
smmsp:x:25:25:SendMail Message Submission Program:/:
listen:x:37:4:Network Admin:/usr/net/nls:
gdm:x:50:50:GDM Reserved UID:/:
websrvd:x:80:80:WebServer Reserved UID:/:
```

```

postgres:x:90:90:PostgreSQL Reserved UID:/:usr/bin/pfksh
svctag:x:95:12:Service Tag UID:/:
nobody:x:60001:60001:NFS Anonymous Access User:/:
noaccess:x:60002:60002:No Access User:/:
nobody4:x:65534:65534:SunOS 4.x NFS Anonymous Access User:/:

```

TABLE 4-6 Default passwd File Entries

User Name	User ID	Description
root	0	Superuser account
daemon	1	Umbrella system daemon associated with routine system tasks
bin	2	Administrative daemon associated with running system binaries to perform some routine system task
sys	3	Administrative daemon associated with system logging or updating files in temporary directories
adm	4	Administrative daemon associated with system logging
lp	71	Line printer daemon
uucp	5	Daemon associated with uucp functions
nuucp	6	Another daemon associated with uucp functions
smmsp	25	Sendmail message submission program daemon
webservd	80	Account reserved for WebServer access
postgres	90	Account reserved for PostgreSQL access
svctag	95	Service Tag Registry access
gdm	50	GNOME Display Manager daemon
listen	37	Network listener daemon
nobody	60001	Account reserved for anonymous NFS access.
noaccess	60002	Assigned to a user or a process that needs access to a system through some application but without actually logging in
nobody4	65534	SunOS 4.0 or 4.1 version of the nobody user account

Fields in the shadow File

The fields in the shadow file are separated by colons and contain the following information:

```
username:password:lastchg:min:max:warn:inactive:expire
```

For example:

```
rimmer:86Kg/MNT/dGu.:8882:0::5:20:8978
```

For a complete description of the fields in the shadow file, see the [shadow\(4\)](#) and [crypt\(1\)](#) man pages.

Fields in the group File

The fields in the group file are separated by colons and contain the following information:

```
group-name:group-password:gid:user-list
```

For example:

```
bin::2:root,bin,daemon
```

For a complete description of the fields in the group file, see the [group\(4\)](#) man page.

Default group File

The default group file contains the following system groups that support some system-wide task, such as printing, network administration, or electronic mail. Many of these groups having corresponding entries in the `passwd` file.

```
root::0:
other::1:root
bin::2:root,daemon
sys::3:root,bin,adm
adm::4:root,daemon
uucp::5:root
mail::6:root
tty::7:root,adm
lp::8:root,adm
nuucp::9:root
staff::10:
daemon::12:root
sysadmin::14:
smsp::25:
gdm::50:
webservd::80:
postgres::90:
nobody::60001:
noaccess::60002:
nogroup::65534:
```


TABLE 4-7 Default group File Entries

Group Name	Group ID	Description
root	0	Superuser group
other	1	Optional group
bin	2	Administrative group associated with running system binaries
sys	3	Administrative group associated with system logging or temporary directories
adm	4	Administrative group associated with system logging
uucp	5	Group associated with uucp functions
mail	6	Electronic mail group
tty	7	Group associated with tty devices
lp	8	Line printer group
nuucp	9	Group associated with uucp functions
staff	10	General administrative group.
daemon	12	Group associated with routine system tasks
sysadmin	14	Administrative group associated with legacy Admintool and Solstice AdminSuite tools
smmsp	25	Daemon for Sendmail message submission program
gdm	50	Group reserved for the GNOME Display Manager daemon
webservd	80	Group reserved for WebServer access
postgres	90	Group reserved for PostgreSQL access
nobody	60001	Group assigned for anonymous NFS access
noaccess	60002	Group assigned to a user or a process that needs access to a system through some application but without actually logging in
nogroup	65534	Group assigned to a user who is not a member of a known group

Tools for Managing User Accounts and Groups

The following table lists the recommended tools for managing users and groups. These tools are included in the Solaris Management Console suite of tools. For information about starting and using the Solaris Management Console, see [Chapter 2, “Working With the Solaris Management Console \(Tasks\)”](#).

TABLE 4–8 Tools for Managing Users and Groups

Solaris Management Tool	Purpose
Users	Manage users accounts
User Templates	Create a set of attributes for a specific kind of user like students, engineers, or instructors
Rights	Manage RBAC rights
Administrative Roles	Manage RBAC administrative roles
Groups	Manage group information
Projects	Manage project information
Mailing Lists	Manage mailing lists

Use the Solaris Management Console online help for information on performing these tasks.

For information on the Solaris commands that can be used to manage user accounts and groups, see [Table 1–5](#). These commands provide the same functionality as the Solaris management tools, including authentication and name service support.

Tasks for Solaris User and Group Management Tools

The Solaris user management tools enable you to manage user accounts and groups on a local system or in a name service environment.

This table describes the tasks that you can perform with the Users tool's User Accounts feature.

TABLE 4–9 Task Descriptions for User Accounts Tool

Task	Description
Add a user.	Adds a user to the local system or name service.
Create a user template.	Creates a template of predefined user attributes for creating users of the same group, such as students, contractors, or engineers.

TABLE 4-9 Task Descriptions for User Accounts Tool (Continued)

Task	Description
Add a user with a user template.	Adds a user with a template so that user attributes are predefined.
Clone a user template.	Clones a user template if you would like to use a similar set of predefined user attributes. Then, change only some of the attributes as needed.
Set up user properties.	Sets up user properties in advance of adding users. Properties include specifying whether a user template is used when adding a user, and whether the home directory or mail box is deleted by default when removing a user.
Add multiple users.	Adds multiple users to the local system or name service by specifying a text file, typing each name, or automatically generating a series of user names.
View or change user properties.	Displays or changes user properties such as login shell, password, or password options.
Assign rights to users.	Assigns RBAC rights to users that will allow them to perform specific administration tasks.
Remove a user.	Removes the user from the local system or the name service. Optionally, you can also specify whether the user's home directory or mailbox is removed. The user is also removed from any groups or roles.

For information about adding a user to the local system or name service, see [“What Are User Accounts and Groups?”](#) on page 86 and [“User Account Components”](#) on page 87.

TABLE 4-10 Task Descriptions for Rights Tool

Task	Description
Grant a right.	Grants a user a right to run a specific command or application that was previously only available to an administrator.
View or change existing rights properties.	Displays or changes existing rights.
Add an authorization.	Adds an authorization, which is a discrete right granted to a role or a user.
View or change an authorization.	Displays or changes existing authorizations.

For more information on granting rights to users, see [“Contents of Rights Profiles”](#) in *System Administration Guide: Security Services*.

TABLE 4-11 Task Descriptions for Administrative Roles Tool

Task	Description
Add an administrative role.	Adds a role that someone would use to perform a specific administrative task.
Assign rights to an administrative role.	Assigns specific rights to a role that enable someone to perform a task.
Change an administrative role.	Adds or removes rights from a role.

For more information on using administrative roles, see [“How to Plan Your RBAC Implementation”](#) in *System Administration Guide: Security Services*.

TABLE 4-12 Task Descriptions for Groups Tool

Task	Description
Add a group.	Adds a group to the local system or name service so that the group name is available before you add the user.
Add a user to a group.	Adds a user to a group if the user needs access to group-owned files.
Remove a user from a group.	Removes a user from a group if the user no longer requires group file access.

For information on adding users to groups, see [“UNIX Groups”](#) on page 89.

TABLE 4-13 Task Descriptions for Mailing Lists Tool

Task	Description
Create a mailing list.	Creates a mailing list, which is a list of user names for sending email messages.
Change a mailing list name.	Changes the mailing list after it is created.
Remove a mailing list.	Removes a mailing list if it is no longer used.

For information on creating mailing lists, see the Solaris Management Console online help.

TABLE 4-14 Task Descriptions for Projects Tool

Task	Description
Create or clone a project.	Creates a new project or clones an existing project if the existing project has attributes similar to what you need for the new project.

TABLE 4-14 Task Descriptions for Projects Tool (Continued)

Task	Description
Modify or view project attributes.	Displays or changes existing project attributes.
Delete a project.	Removes a project if the project is no longer used.

Managing Users and Resources With Projects

Users and groups can be members of a *project*, an identifier that indicates a workload component that can be used as the basis of system usage or resource allocation charge-back. Projects are part of the Solaris resource management feature that is used to manage system resources.

Users need to be a member of a project to successfully log in to a system running the Solaris 9 release. By default, users are a member of the `group.staff` project when the Solaris 9 release is installed and no other project information is configured.

User project information is stored in the `/etc/project` file, which can be stored on the local system (files), the NIS name service, or the LDAP directory service. You can use the Solaris Management Console to manage project information.

The `/etc/project` file must exist for users to log in successfully, but requires no administration if you are not using projects.

For more information on using or setting up projects, see [Chapter 2, “Projects and Tasks \(Overview\)”](#), in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*.

Customizing a User's Work Environment

Part of setting up a user's home directory is providing user initialization files for the user's login shell. A *user initialization file* is a shell script that sets up a work environment for a user after the user logs in to a system. Basically, you can perform any task in a user initialization file that you can do in a shell script. However, a user initialization file's primary job is to define the characteristics of a user's work environment, such as a user's search path, environment variables, and windowing environment. Each login shell has its own user initialization file or files, which are listed in the following table.

TABLE 4-15 User Initialization Files for Bourne, C, and Korn Shells

Shell	User Initialization File	Purpose
Bourne	<code>\$HOME/.profile</code>	Defines the user's environment at login

TABLE 4-15 User Initialization Files for Bourne, C, and Korn Shells (Continued)

Shell	User Initialization File	Purpose
C	<code>\$HOME/.cshrc</code>	Defines the user's environment for all C shells and is invoked after login shell
	<code>\$HOME/.login</code>	Defines the user's environment at login
Korn	<code>\$HOME/.profile</code>	Defines the user's environment at login
	<code>\$HOME/\$ENV</code>	Defines user's environment at login in the file and is specified by the Korn shell's ENV environment variable

TABLE 4-16 Default User Initialization Files

Shell	Default File
C	<code>/etc/skel/local.login</code>
	<code>/etc/skel/local.cshrc</code>
Bourne or Korn	<code>/etc/skel/local.profile</code>

You can use these files as a starting point and then modify them to create a standard set of files that provide the work environment common to all users. You can also modify these files to provide the working environment for different types of users. Although you cannot create customized user initialization files with the Users tool, you can populate a user's home directory with user initialization files that are located in a specified "skeleton" directory. You can accomplish this task by creating a user template with the User Templates tool and then specifying a skeleton directory from which to copy user initialization files.

For step-by-step instructions on how to create sets of user initialization files for different types of users, see ["How to Customize User Initialization Files"](#) on page 115.

When you use the Users tool to create a new user account and select the create home directory option, the following files are created, depending on which login shell is selected.

TABLE 4-17 Files Created by Users Tool When Adding a User

Shell	Files Created
C	The <code>/etc/skel/local.cshrc</code> and the <code>/etc/skel/local.login</code> files are copied into the user's home directory and are renamed <code>.cshrc</code> and <code>.login</code> , respectively.
Bourne and Korn	The <code>/etc/skel/local.profile</code> file is copied into the user's home directory and renamed <code>.profile</code> .

Customizing the Bash Shell

To customize your bash shell, add the information to the `.bashrc` file that is located in your home directory. The initial user that is created when you install Oracle Solaris has a `.bashrc` file that sets the `PATH`, `MANPATH`, and command prompt. For more information, see the `bash(1)` man page.

Using Site Initialization Files

The user initialization files can be customized by both the administrator and the user. This important task can be accomplished with centrally located and globally distributed user initialization files that are called, *site initialization files*. Site initialization files enable you to continually introduce new functionality to the user's work environment, while enabling the user to customize the user's initialization file.

When you reference a site initialization file in a user initialization file, all updates to the site initialization file are automatically reflected when the user logs in to the system or when a user starts a new shell. Site initialization files are designed for you to distribute site-wide changes to users' work environments that you did not anticipate when you added the users.

You can customize a site initialization file the same way that you customize a user initialization file. These files typically reside on a server, or set of servers, and appear as the first statement in a user initialization file. Also, each site initialization file must be the same type of shell script as the user initialization file that references it.

To reference a site initialization file in a C-shell user initialization file, place a line at the beginning of the user initialization file similar to the following line:

```
source /net/machine-name/export/site-files/site-init-file
```

To reference a site initialization file in a Bourne-shell or Korn-shell user initialization file, place a line at the beginning of the user initialization file similar to the following line:

```
./net/machine-name/export/site-files/site-init-file
```

Avoiding Local System References

Do not add specific references to the local system in the user initialization file. The instructions in a user initialization file should be valid, regardless of which system the user logs into.

For example:

- To make a user's home directory available anywhere on the network, always refer to the home directory with the variable `$HOME`. For example, use `$HOME/bin` instead of `/export/home/username/bin`. The `$HOME` variable works when the user logs in to another system, and the home directories are auto-mounted.
- To access files on a local disk, use global path names, such as `/net/system-name/directory-name`. Any directory referenced by `/net/system-name` can be mounted automatically on any system on which the user logs in, assuming the system is running AutoFS.

Shell Features

The following table lists basic shell features that each shell provides, which can help you determine what you can and can't do when creating user initialization files for each shell.

TABLE 4-18 Basic Features of Bourne, C, and Korn Shells

Feature	Bourne	C	Korn
Known as the standard shell in UNIX	Applicable	No applicable	Not applicable
Compatible syntax with Bourne shell	-	Not applicable	Applicable
Job control	Applicable	Applicable	Applicable
History list	Not applicable	Applicable	Applicable
Command-line editing	Not applicable	Applicable	Applicable
Aliases	Not applicable	Applicable	Applicable
Single-character abbreviation for login directory	Not applicable	Applicable	Applicable
Protection from overwriting (<code>noclobber</code>)	Not applicable	Applicable	Applicable
Setting to ignore Control-D (<code>ignoreeof</code>)	Not applicable	Applicable	Applicable
Enhanced <code>cd</code> command	Not applicable	Applicable	Applicable
Initialization file separate from <code>.profile</code>	Not applicable	Applicable	Applicable
Logout file	Not applicable	Applicable	N/A

Shell Environment

A shell maintains an environment that includes a set of variables defined by the `login` program, the system initialization file, and the user initialization files. In addition, some variables are defined by default.

A shell can have two types of variables:

- **Environment variables** – Variables that are exported to all processes spawned by the shell. Their settings can be seen with the `env` command. A subset of environment variables, such as `PATH`, affects the behavior of the shell itself.
- **Shell (local) variables** – Variables that affect only the current shell. In the C shell, a set of these shell variables have a special relationship to a corresponding set of environment variables. These shell variables are `user`, `term`, `home`, and `path`. The value of the environment variable counterpart is initially used to set the shell variable.

In the C shell, you use lowercase names with the `set` command to set shell variables. You use uppercase names with the `setenv` command to set environment variables. If you set a shell variable, the shell sets the corresponding environment variable. Likewise, if you set an environment variable, the corresponding shell variable is also updated. For example, if you update the `path` shell variable with a new path, the shell also updates the `PATH` environment variable with the new path.

In the Bourne and Korn shells, you can use the uppercase variable name that is equal to some value to set both shell and environment variables. Also, use the `export` command to activate the variables for any subsequently executed commands.

For all shells, you generally refer to shell and environment variables by their uppercase names.

In a user initialization file, you can customize a user's shell environment by changing the values of the predefined variables or by specifying additional variables. The following table shows how to set environment variables in a user initialization file.

TABLE 4-19 Setting Environment Variables in a User Initialization File

Shell Type	Line to Add to the User Initialization File
C shell	<code>setenv VARIABLE value</code> Example: <code>setenv MAIL /var/mail/ripley</code>
Bourne or Korn shell	<code>VARIABLE=value; export VARIABLE</code> Example: <code>MAIL=/var/mail/ripley;export MAIL</code>

The following table describes environment variables and shell variables that you might want to customize in a user initialization file. For more information about variables that are used by the different shells, see the [sh\(1\)](#), [ksh\(1\)](#), or [csh\(1\)](#) man pages.

TABLE 4-20 Shell and Environment Variable Descriptions

Variable	Description
CDPATH, or <code>cdpath</code> in the C shell	Sets a variable used by the <code>cd</code> command. If the target directory of the <code>cd</code> command is specified as a relative path name, the <code>cd</code> command first looks for the target directory in the current directory (<code>.</code>). If the target is not found, the path names listed in the CDPATH variable are searched consecutively until the target directory is found and the directory change is completed. If the target directory is not found, the current working directory is left unmodified. For example, the CDPATH variable is set to <code>/home/jean</code> , and two directories exist under <code>/home/jean</code> , <code>bin</code> , and <code>rje</code> . If you are in the <code>/home/jean/bin</code> directory and type <code>cd rje</code> , you change directories to <code>/home/jean/rje</code> , even though you do not specify a full path.
<code>history</code>	Sets the history for the C shell.
HOME, or <code>home</code> in the C shell	Sets the path to the user's home directory.
LANG	Sets the locale.
LOGNAME	Defines the name of the user currently logged in. The default value of LOGNAME is set automatically by the <code>login</code> program to the user name specified in the <code>passwd</code> file. You should only need to refer to, not reset, this variable.
LPDEST	Sets the user's default printer.
MAIL	Sets the path to the user's mailbox.
MANPATH	Sets the hierarchies of man pages that are available.
PATH, or <code>path</code> in the C shell	Specifies, in order, the directories that the shell searches to find the program to run when the user types a command. If the directory is not in the search path, users must type the complete path name of a command. As part of the <code>login</code> process, the default PATH is automatically defined and set as specified in <code>.profile</code> (Bourne or Korn shell) or <code>.cshrc</code> (C shell). The order of the search path is important. When identical commands exist in different locations, the first command found with that name is used. For example, suppose that PATH is defined in Bourne and Korn shell syntax as <code>PATH=/bin:/usr/bin:/usr/sbin:\$HOME/bin</code> and a file named <code>sample</code> resides in both <code>/usr/bin</code> and <code>/home/jean/bin</code> . If the user types the command <code>sample</code> without specifying its full path name, the version found in <code>/usr/bin</code> is used.
<code>prompt</code>	Defines the shell prompt for the C shell.
PS1	Defines the shell prompt for the Bourne or Korn shell.

TABLE 4–20 Shell and Environment Variable Descriptions (Continued)

Variable	Description
SHELL, or shell in the C shell	Sets the default shell used by make, vi, and other tools.
TERMINFO	Names a directory where an alternate terminfo database is stored. Use the TERMINFO variable in either the <code>/etc/profile</code> or <code>/etc/.login</code> file. For more information, see the <code>terminfo(4)</code> man page. When the TERMINFO environment variable is set, the system first checks the TERMINFO path defined by the user. If the system does not find a definition for a terminal in the TERMINFO directory defined by the user, it searches the default directory, <code>/usr/share/lib/terminfo</code> , for a definition. If the system does not find a definition in either location, the terminal is identified as “dumb.”
TERM, or term in the C shell	Defines the terminal. This variable should be reset in either the <code>/etc/profile</code> or <code>/etc/.login</code> file. When the user invokes an editor, the system looks for a file with the same name that is defined in this environment variable. The system searches the directory referenced by TERMINFO to determine the terminal characteristics.
TZ	Sets the time zone. The time zone is used to display dates, for example, in the <code>ls -l</code> command. If TZ is not set in the user's environment, the system setting is used. Otherwise, Greenwich Mean Time is used.

The PATH Variable

When the user executes a command by using the full path, the shell uses that path to find the command. However, when users specify only a command name, the shell searches the directories for the command in the order specified by the PATH variable. If the command is found in one of the directories, the shell executes the command.

A default path is set by the system. However, most users modify it to add other command directories. Many user problems related to setting up the environment and accessing the correct version of a command or a tool can be traced to incorrectly defined paths.

Setting Path Guidelines

Here are some guidelines for setting up efficient PATH variables:

- If you must include the current directory (`.`) in your path, it should be placed last. Including the current directory (`.`) in your path is a security risk because some malicious person could hide a compromised script or executable in the current directory. Consider using absolute path names instead.
- Keep the search path as short as possible. The shell searches each directory in the path. If a command is not found, long searches can slow down system performance.
- The search path is read from left to right, so you should put directories for commonly used commands at the beginning of the path.

- Make sure that directories are not duplicated in the path.
- Avoid searching large directories, if possible. Put large directories at the end of the path.
- Put local directories before NFS mounted directories to lessen the chance of “hanging” when the NFS server does not respond. This strategy also reduces unnecessary network traffic.

Locale Variables

The LANG and LC environment variables specify the locale-specific conversions and conventions for the shell. These conversions and conventions include time zones, collation orders, and formats of dates, time, currency, and numbers. In addition, you can use the `stty` command in a user initialization file to indicate whether the terminal session will support multibyte characters.

The LANG variable sets all possible conversions and conventions for the given locale. You can set various aspects of localization separately through these LC variables: LC_COLLATE, LC_CTYPE, LC_MESSAGES, LC_NUMERIC, LC_MONETARY, and LC_TIME.

The following table describes some of the values for the LANG and LC environment variables.

TABLE 4-21 Values for LANG and LC Variables

Value	Locale
de_DE.ISO8859-1	German
en_US.UTF-8	American English (UTF-8)
es_ES.ISO8859-1	Spanish
fr_FR.ISO8859-1	French
it_IT.ISO8859-1	Italian
ja_JP.eucJP	Japanese (EUC)
ko_KR.EUC	Korean (EUC)
sv_SE.ISO8859-1	Swedish
zh_CN.EUC	Simplified Chinese (EUC)
zh_TW.EUC	Traditional Chinese (EUC)

For more information on supported locales, see the *International Language Environments Guide*.

EXAMPLE 4-1 Setting the Locale Using the LANG Variables

The following examples show how to set the locale by using the LANG environment variables. In a C-shell user initialization file, you would add the following:

```
setenv LANG de_DE.ISO8859-1
```

In a Bourne-shell or Korn-shell user initialization file, you would add the following:

```
LANG=de_DE.ISO8859-1; export LANG
```

Default File Permissions (umask)

When you create a file or directory, the default file permissions assigned to the file or directory are controlled by the *user mask*. The user mask is set by the `umask` command in a user initialization file. You can display the current value of the user mask by typing `umask` and pressing Return.

The user mask contains the following octal values:

- The first digit sets permissions for the user
- The second digit sets permissions for group
- The third digit sets permissions for other, also referred to as `world`

Note that if the first digit is zero, it is not displayed. For example, if the user mask is set to 022, 22 is displayed.

To determine the `umask` value that you want to set, subtract the value of the permissions you want from 666 (for a file) or 777 (for a directory). The remainder is the value to use with the `umask` command. For example, suppose you want to change the default mode for files to 644 (`rw-r--r--`). The difference between 666 and 644 is 022, which is the value you would use as an argument to the `umask` command.

You can also determine the `umask` value you want to set by using the following table. This table shows the file and directory permissions that are created for each of the octal values of `umask`.

TABLE 4-22 Permissions for `umask` Values

umask Octal Value	File Permissions	Directory Permissions
0	rw-	rwx
1	rw-	rw-
2	r--	r-x
3	r--	r--
4	-w-	-wx

TABLE 4-22 Permissions for umask Values (Continued)

umask Octal Value	File Permissions	Directory Permissions
5	-w-	-w-
6	--x	--x
7	--- (none)	--- (none)

The following line in a user initialization file sets the default file permissions to `rw-rw-rw-`.

```
umask 000
```

User and Site Initialization Files Examples

The following sections provide examples of user and site initialization files that you can use to start customizing your own initialization files. These examples use system names and paths that you need to change for your particular site.

EXAMPLE 4-2 The .profile File

```
(Line 1) PATH=$PATH:$HOME/bin:/usr/local/bin:/usr/ccs/bin:.
(Line 2) MAIL=/var/mail/$LOGNAME
(Line 3) NNTPSERVER=server1
(Line 4) MANPATH=/usr/share/man:/usr/local/man
(Line 5) PRINTER=printer1
(Line 6) umask 022
(Line 7) export PATH MAIL NNTPSERVER MANPATH PRINTER
```

1. Defines the user's shell search path.
2. Defines the path to the user's mail file.
3. Defines the user's Usenet news server.
4. Defines the user's search path for man pages.
5. Defines the user's default printer.
6. Sets the user's default file creation permissions.
7. Sets the listed environment variables.

EXAMPLE 4-3 The .cshrc File

```
(Line 1) set path=( $PATH $HOME/bin /usr/local/bin /usr/ccs/bin)
(Line 2) setenv MAIL /var/mail/$LOGNAME
(Line 3) setenv NNTPSERVER server1
(Line 4) setenv PRINTER printer1
(Line 5) alias h history
(Line 6) umask 022
(Line 7) source /net/server2/site-init-files/site.login
```

1. Defines the user's shell search path.

EXAMPLE 4-3 The `.cshrc` File (Continued)

2. Defines the path to the user's mail file.
3. Defines the user's Usenet news server.
4. Defines the user's default printer.
5. Creates an alias for the `history` command. The user needs to type only `h` to run the `history` command.
6. Sets the user's default file creation permissions.
7. Sources the site initialization file.

EXAMPLE 4-4 Site Initialization File

The following shows an example site initialization file in which a user can choose a particular version of an application.

```
# @(#)site.login
main:
echo "Application Environment Selection"
echo ""
echo "1. Application, Version 1"
echo "2. Application, Version 2"
echo ""
echo -n "Type 1 or 2 and press Return to set your
application environment: "

set choice = $<

if ( $choice !~ [1-2] ) then
goto main
endif

switch ($choice)

case "1":
setenv APPHOME /opt/app-v.1
breaksw

case "2":
setenv APPHOME /opt/app-v.2
endsw
```

This site initialization file could be referenced in a user's `.cshrc` file (C shell users only) with the following line:

```
source /net/server2/site-init-files/site.login
```

In this line, the site initialization file is named `site.login` and is located on a server named `server2`. This line also assumes that the automounter is running on the user's system.

Managing User Accounts and Groups (Tasks)

This chapter describes how to set up and maintain user accounts and groups.

For information on the procedures associated with setting up and maintaining user accounts and groups, see the following:

- [“Setting Up User Accounts \(Task Map\)” on page 113](#)
- [“Maintaining User Accounts \(Task Map\)” on page 123](#)

For information on the procedures associated with setting up and maintaining user accounts and groups, see [“Setting Up User Accounts \(Task Map\)” on page 113](#).

For background information about managing user accounts and groups, see [Chapter 4, “Managing User Accounts and Groups \(Overview\)”](#).

Setting Up User Accounts (Task Map)

Task	Description	For Instructions
Gather user information.	Use a standard form to gather user information to help you keep user information organized.	“Gathering User Information” on page 114
Customize user initialization files.	You can set up user initialization files (.cshrc, .profile, .login), so that you can provide new users with consistent environments.	“How to Customize User Initialization Files” on page 115

Task	Description	For Instructions
Add a group.	Use Oracle Solaris command-line interface tools to add a group.	“How to Add a Group With the Solaris Management Console’s Groups Tool” on page 116 “Adding Groups and Users With Command-Line Tools” on page 119
Add a user.	You can add a user with the following tools: Solaris Management Console’s Users tool Solaris command-line interface tools	“How to Add a User With the Solaris Management Console’s Users Tool” on page 117 “Adding Groups and Users With Command-Line Tools” on page 119
Set up a user template.	You can create a user template, so that you don’t have to manually add all similar user properties.	See Solaris Management Console online help
Add rights or a role to a user.	You can add rights or a role to a user so that the user can perform a specific command or task.	See Solaris Management Console online help
Share the user’s home directory.	You must share the user’s home directory so that the directory can be remotely mounted from the user’s system.	“How to Share a User’s Home Directory” on page 120
Mount the user’s home directory.	You must mount the user’s home directory on the user’s system.	“How to Mount a User’s Home Directory” on page 121

Gathering User Information

You can create a form such as the following to gather information about users before adding their accounts.

Item	Description
User Name:	
Role Name:	
Profiles or Authorizations:	
UID:	
Primary Group:	
Secondary Groups:	

Item	Description
Comment:	
Default Shell:	
Password Status and Aging:	
Home Directory Path Name:	
Mounting Method:	
Permissions on Home Directory:	
Mail Server:	
Department Name:	
Department Administrator:	
Manager:	
Employee Name:	
Employee Title:	
Employee Status:	
Employee Number:	
Start Date:	
Add to These Mail Aliases:	
Desktop System Name:	

▼ How to Customize User Initialization Files

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Create a skeleton directory for each type of user.

```
# mkdir /shared-dir/skel/user-type
```

shared-dir The name of a directory that is available to other systems on the network.

user-type The name of a directory to store initialization files for a type of user.

- 3 **Copy the default user initialization files into the directories that you created for different types of users.**

```
# cp /etc/skel/local.cshrc /shared-dir/skel/user-type/.cshrc
# cp /etc/skel/local.login /shared-dir/skel/user-type/.login
# cp /etc/skel/local.profile /shared-dir/skel/user-type/.profile
```

Note – If the account has profiles assigned to it, then the user has to launch a special version of the shell called a profile shell to use commands (with any security attributes) that are assigned to the profile. There are three *profile shells* corresponding to the types of shells: `pfsh` (Bourne shell), `pcfsh` (C shell), and `pfksh` (Korn shell). For information about profile shells, see [“Role-Based Access Control \(Overview\)”](#) in *System Administration Guide: Security Services*.

- 4 **Edit the user initialization files for each user type and customize them based on your site's needs.**

For a detailed description on the ways to customize the user initialization files, see [“Customizing a User's Work Environment”](#) on page 101.

- 5 **Set the permissions for the user initialization files.**

```
# chmod 744 /shared-dir/skel/user-type/*
```

- 6 **Verify that the permissions for the user initialization files are correct.**

```
# ls -la /shared-dir/skel/*
```

Example 5–1 Customizing User Initialization Files

The following example shows how to customize the C-shell user initialization file in the `/export/skel/enduser` directory designated for a particular type of user. For an example of a `.cshrc` file, see [Example 4–3](#).

```
# mkdir /export/skel/enduser
# cp /etc/skel/local.cshrc /export/skel/enduser/.cshrc
```

```
(Edit .cshrc file)
# chmod 744 /export/skel/enduser/*
```

▼ How to Add a Group With the Solaris Management Console's Groups Tool

You can add existing users to the group when you add the group. Or, you can just add the group and then add the user to the group when you add the user.

- 1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.
- 2 Start the Solaris Management Console.**

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see “How to Start the Console as Superuser or as a Role” on page 44 or “How to Start the Solaris Management Console in a Name Service Environment” on page 51.
- 3 Click the This Computer icon under the Management Tools icon in the Navigation pane.**

A list of categories is displayed.
- 4 (Optional) Select the appropriate toolbox for your name service environment.**
- 5 Click the System Configuration icon.**
- 6 Click the User icon and provide the superuser password or the role password.**
- 7 Click the Groups icon. Select Add Group from the Action menu.**

Use the Context help to add a group to the system.
- 8 Identify the group name at the Group Name prompt under Group Identification.**

For example, `mechanoids`.
- 9 Identify the group number at the Group ID number prompt.**

For example, `GID 101`.
- 10 Click OK.**

▼ How to Add a User With the Solaris Management Console's Users Tool

Use the following procedure to add a user with the Solaris Management Console's Users tool.

- 1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Start the Solaris Management Console.

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role”](#) on page 44 or [“How to Start the Solaris Management Console in a Name Service Environment”](#) on page 51.

3 Click the This Computer icon under the Management Tools icon in the Navigation pane.

A list of categories is displayed.

4 (Optional) Select the appropriate toolbox for your name service environment.

5 Click the System Configuration icon.

6 Click the User icon and provide the superuser password or the role password.

7 Click the User Accounts icon.

Use the Context help to add a user to the system.

8 Select Add User → With Wizard from the Action menu.

Click Next between the steps below.

a. Identify the user name or login name at the User Name prompt.

For example, kryten

b. (Optional) Identify the user's full name at the Full Name prompt.

For example, kryten series 3000.

c. (Optional) Provide a further description of this user at the Description prompt.

d. Provide the user ID at the User ID Number prompt.

For example, 1001.

e. Select the User Must Use This Password At First Login option.

f. At the Password prompt, provide a password for the user.

g. At the Confirm Password prompt, confirm the password.

h. Select the user's primary group.

For example, mechanoids.

- i. Create the user's home directory by accepting the defaults at the Server and Path prompts.
- j. Specify the mail server.
- k. Review the information that you provided, and if necessary, go back and correct any wrong information.
- l. Click Finish.

Adding Groups and Users With Command-Line Tools

This section provides examples of adding users and groups with command-line tools.

Adding a Group and User With the `groupadd` and `useradd` Commands

The following example shows how to use the `groupadd` and `useradd` commands to add the group `scutters` and the user `scutter1` to files on the local system. These commands cannot be used to manage users in a name service environment.

```
# groupadd -g 102 scutters
# useradd -u 1003 -g 102 -d /export/home/scutter1 -s /bin/csh \
-c "Scutter 1" -m -k /etc/skel scutter1
64 blocks
```

For more information, see the [groupadd\(1M\)](#) and [useradd\(1M\)](#) man pages.

Adding a Group and User With the `smgroup` and `smuser` Commands

The following example shows how to use the `smgroup` and `smuser` commands to add the group `gelfs` and the user `camille` to the NIS domain `solar.com` on the host `starlite`.

```
# /usr/sadm/bin/smgroup add -D nis:/starlitesolar.com -- -g 103 -n gelfs
# /usr/sadm/bin/smuser add -D nis:/starlite/solar.com -- -u 1004
-n camille -c "Camille G." -d /export/home/camille -s /bin/csh -g gelfs
```

For more information, see the [smgroup\(1M\)](#) and [smuser\(1M\)](#) man pages.

Setting Up Home Directories With the Solaris Management Console

Keep the following in mind when using Solaris Management Console to manage user home directories:

- If you use the Users tool's Add User Wizard to add a user account and you specify the user's home directory as `/export/home/username`, the home directory is automatically set up to automount. Also, the following entry is added to the `passwd` file.

`/home/username`
- There is only way you can use Users tool to set up a user account that does not automount the home directory. First, set up a user account template that disables this feature. Then, add users with this template. You cannot disable this feature with the Add User Wizard.
- You can use the `smuser add` command with the `-x autohome=N` option to add a user without automounting the user's home directory. However, there is no option to the `smuser delete` command to remove the home directory after the user is added. You would have to remove the user and the user's home directory with the Users tool.

▼ How to Share a User's Home Directory

Use the following procedure to share a user's home directory.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Verify that the `mountd` daemon is running.

In this release, `mountd` is now started as part of the NFS server service. To see if the `mountd` daemon is running, type the following command:

```
# svcs network/nfs/server
STATE          STIME      FMRI
online         Aug_26    svc:/network/nfs/server:default
```

3 If the `mountd` daemon is not running, start it.

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

4 List the file systems that are shared on the system.

```
# share
```


- 5 Select one of the following based on whether the file system that contains the user's home directory is already shared.
 - a. If the user's home directory is already shared, go to the step 8.
 - b. If the user's home directory is not shared, go to [Step 6](#).
- 6 Edit the `/etc/dfs/dfstab` file and add the following line:


```
share -F nfs /file-system
```

/file-system is the file system that contains the user's home directory that you need to share. By convention, the file system is `/export/home`.
- 7 Share the file systems listed in the `/etc/dfs/dfstab` file.


```
# shareall -F nfs
```

This command executes all the share commands in the `/etc/dfs/dfstab` file so that you do not have to wait to reboot the system.
- 8 Verify that a user's home directory is shared.


```
# share
```

Example 5-2 Sharing a User's Home Directory

The following example shows how to share the `/export/home` directory.

```
# svcs network/nfs/server
# svcadm network/nfs/server
# share
# vi /etc/dfs/dfstab

(The line share -F nfs /export/home is added.)
# shareall -F nfs
# share
-           /usr/dist           ro   ""
-           /export/home/user-name   rw   ""
```

See Also If the user's home directory is not located on the user's system, you have to mount the user's home directory from the system where it is located. For detailed instructions, see [“How to Mount a User's Home Directory”](#) on page 121.

▼ How to Mount a User's Home Directory

For information on automounting a home directory, see [“Task Overview for Autofs Administration”](#) in *System Administration Guide: Network Services*.

1 Make sure that the user's home directory is shared.

For more information, see [“How to Share a User's Home Directory”](#) on page 120.

2 Log in as superuser on the user's system.**3 Edit the `/etc/vfstab` file and create an entry for the user's home directory.**

```
system-name:/export/home/user-name - /export/home/username nfs - yes rw
```

`system-name` The name of the system where the home directory is located.

`/export/home/username` The name of the user's home directory that will be shared. By convention, `/export/home/username` contains user home directories. However, you can use a different file system.

`-` Required placeholders in the entry.

`/export/home/username` The name of the directory where the user's home directory will be mounted.

For more information about adding an entry to the `/etc/vfstab` file, see [“Overview of Mounting and Unmounting File Systems”](#) in *System Administration Guide: Devices and File Systems*.

4 Create the mount point for the user's home directory.

```
# mkdir -p /export/home/username
```

5 Mount the user's home directory.

```
# mountall
```

All entries in the current `vfstab` file (whose mount at boot fields are set to yes) are mounted.

6 Verify that the home directory is mounted.

```
# mount | grep username
```

Example 5-3 Mounting a User's Home Directory

The following example shows how to mount user `ripley`'s home directory.

```
# vi /etc/vfstab
```

```
(The line venus:/export/home/ripley - /export/home/ripley
```

```
nfs - yes rw is added.)
```

```
# mkdir -p /export/home/ripley
```

```
# mountall
```

```
# mount
```

```
/ on /dev/dsk/c0t0d0s0 read/write/setuid/intr/largefiles/xattr/onerror=panic/dev=...
```

```
/devices on /devices read/write/setuid/dev=46c0000 on Thu Jan 8 09:38:19 2004
```

```
/usr on /dev/dsk/c0t0d0s6 read/write/setuid/intr/largefiles/xattr/onerror=panic/dev=...
```

```
/proc on /proc read/write/setuid/dev=4700000 on Thu Jan 8 09:38:27 2004
```

```

/etc/mnttab on mnttab read/write/setuid/dev=47c0000 on Thu Jan  8 09:38:27 2004
/dev/fd on fd read/write/setuid/dev=4800000 on Thu Jan  8 09:38:30 2004
/var/run on swap read/write/setuid/xattr/dev=1 on Thu Jan  8 09:38:30 2004
/tmp on swap read/write/setuid/xattr/dev=2 on Thu Jan  8 09:38:30 2004
/export/home on /dev/dsk/c0t0d0s7 read/write/setuid/intr/largefiles/xattr/onerror=...
/export/home/ripley on venus:/export/home/ripley remote/read/write/setuid/xattr/dev=...

```

Maintaining User Accounts (Task Map)

Task	Description	Instructions
Modify a group.	You can modify a group's name or the users in a group by using the Groups tool.	“How to Modify a Group” on page 125
Delete a group.	You can delete a group if it is no longer needed.	“How to Delete a Group” on page 125
Modify a user account.	<p>Disable a user account:</p> <p>You can temporarily disable a user account if it will be needed in the future.</p> <p>Change a user's password:</p> <p>You might need to change a user's password if the user forgets it.</p> <p>Set password aging:</p> <p>You can force users to change their passwords periodically with User Account tool's Password Options menu.</p>	<p>“How to Disable a User Account” on page 127</p> <p>“How to Change a User's Password” on page 128</p> <p>“How to Set Password Aging on a User Account” on page 129</p>
Delete a user account.	You can delete a user account if it is no longer needed.	“How to Delete a User Account” on page 129

Modifying User Accounts

Unless you define a user name or UID number that conflicts with an existing one, you should never need to modify a user account's user name or UID number.

Use the following steps if two user accounts have duplicate user names or UID numbers:

- If two user accounts have duplicate UID numbers, use the Users tool to remove one account and add it again with a different UID number. You cannot use the Users tool to modify a UID number of an existing user account.
- If two user accounts have duplicate user names, use the Users tool to modify one of the accounts and change the user name.

If you do use the Users tool to change a user name, the home directory's ownership is changed, if a home directory exists for the user.

One part of a user account that you can change is a user's group memberships. To add or delete a user's secondary groups, select the Properties option from the Users tool's Action menu. Alternatively, you can use the Groups tool to directly modify a group's member list.

You can also modify the following parts of a user account:

- Description (comment)
- Login shell
- Passwords and password options
- Home directory and home directory access
- Rights and roles

Disabling User Accounts

Occasionally, you might need to temporarily or permanently disable a user account. Disabling or locking a user account means that an invalid password, *LK*, is assigned to the user account, preventing future logins.

The easiest way to disable a user account is to lock the password for an account with Users tool.

You can also enter an expiration date in the account availability section of the User Properties screen. An expiration date enables you to set a limit on how long the account is active.

Other ways to disable a user account: set up password aging or change the user's password.

Deleting User Accounts

When you delete a user account with the Users tool, the software deletes the entries in the passwd and group files. In addition, the files in the user's home directory and mail directory are deleted also.

▼ How to Modify a Group

Use the following procedure to modify a group.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Start the Solaris Management Console.

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see “How to Start the Console as Superuser or as a Role” on page 44 or “How to Start the Solaris Management Console in a Name Service Environment” on page 51.

3 Click the This Computer icon under the Management Tools icon in the Navigation pane.

A list of categories is displayed.

4 (Optional) Select the appropriate toolbox for your name service environment.

5 Click the System Configuration icon.

6 Click the User icon.

7 Provide the superuser password or the role password.

8 Click the Groups icon.

9 Select the group to modify.

For example, select `scutters`.

10 Modify the selected group in the Group Name: text box. Click OK when you are finished.

For example, change `scutters` to `scutter`.

All the users that were in the `scutters` group are now in the `scutter` group.

▼ How to Delete a Group

Use the following procedure to delete a group.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Start the Solaris Management Console.

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role”](#) on page 44 or [“How to Start the Solaris Management Console in a Name Service Environment”](#) on page 51.

3 Click the This Computer icon under the Management Tools icon in the Navigation pane.

A list of categories is displayed.

4 (Optional) Select the appropriate toolbox for your name service environment.

5 Click the System Configuration icon.

6 Click the User icon.

7 Provide the superuser password or the role password.

8 Click the Groups icon.

9 Select the group to delete.

For example, select scutter.

10 Click OK in the pop-up window.

The group is removed from all the users who were a member of this group.

Administering Passwords

You can use the Users tool for password administration. This tool includes the following capabilities:

- Specifying a normal password for a user account
- Enabling users to create their own passwords during their first login
- Disabling or locking a user account
- Specifying expiration dates and password aging information

Note – Password aging is not supported by the NIS name service.

Using Password Aging

If you are using NIS + or the `/etc` files to store user account information, you can set up password aging on a user's password. Starting in the Solaris 9 12/02 release, password aging is also supported in the LDAP directory service.

Password aging enables you to force users to change their passwords periodically or to prevent a user from changing a password before a specified interval. If you want to prevent an intruder from gaining undetected access to the system by using an old and inactive account, you can also set a password expiration date when the account becomes disabled. You can set password aging attributes with the `passwd` command or the Solaris Management Console's Users tool.

For information about starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role” on page 44](#).

▼ How to Disable a User Account

Use the following procedure if you need to disable a user account.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Start the Solaris Management Console.

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role” on page 44](#) or [“How to Start the Solaris Management Console in a Name Service Environment” on page 51](#).

3 Click the This Computer icon under the Management Tools icon in the Navigation pane.

A list of categories is displayed.

4 (Optional) Select the appropriate toolbox for your name service environment.

5 Click the System Configuration icon.

6 Click the User icon and provide the superuser password or the role password.

7 Click the User Accounts icon.

8 Double-click the user.

For example, select `scutter2`.

- 9 **Select the Account is Locked option in the Account Availability section of the General tab features.**
- 10 **Click OK.**

▼ **How to Change a User's Password**

Use the following procedure when a user forgets her password.

- 1 **Become superuser or assume an equivalent role.**
Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.
- 2 **Start the Solaris Management Console.**

```
# /usr/sadm/bin/smc &
```


For more information on starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role”](#) on page 44 or [“How to Start the Solaris Management Console in a Name Service Environment”](#) on page 51.
- 3 **Click the This Computer icon under the Management Tools icon in the Navigation pane.**
A list of categories is displayed.
- 4 **(Optional) Select the appropriate toolbox for your name service environment.**
- 5 **Click the System Configuration icon.**
- 6 **Click the User icon.**
- 7 **Provide the superuser password or the role password.**
- 8 **Click the User Accounts icon, then double-click the user who needs a new password.**
For example, select `scutter1`.
- 9 **Select the Password tab, then select the User Must Use This Password at Next Login option.**
- 10 **Enter the user's new password and click OK.**

▼ How to Set Password Aging on a User Account

Use the following procedure to set password aging on a user account.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Start the Solaris Management Console.

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see “How to Start the Console as Superuser or as a Role” on page 44 or “How to Start the Solaris Management Console in a Name Service Environment” on page 51.

3 Click the This Computer icon under the Management Tools icon in the Navigation pane.

A list of categories is displayed.

4 (Optional) Select the appropriate toolbox for your name service environment.

5 Click the System Configuration icon.

6 Click the User Accounts icon and provide the superuser password or the role password.

7 Click the User Accounts icon.

8 Double-click the user, then select the Password Options tab.

For example, select `scutter2`.

9 Select the Password Options tab.

10 Select the appropriate Password Options in Days option and click OK.

For example, select Users Must Change Within to set a date when the user must change his or her password.

▼ How to Delete a User Account

Use the following procedure to remove a user account.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Start the Solaris Management Console.

```
# /usr/sadm/bin/smc &
```

For more information on starting the Solaris Management Console, see [“How to Start the Console as Superuser or as a Role”](#) on page 44 or [“How to Start the Solaris Management Console in a Name Service Environment”](#) on page 51.

3 Click the This Computer icon under the Management Tools icon in the Navigation pane.

A list of categories is displayed.

4 (Optional) Select the appropriate toolbox for your name service environment.

5 Click the System Configuration icon.

6 Click the User icon.

7 Provide the superuser password or the role password.

8 Click the User Accounts icon.

9 Double-click the user account to be removed.

For example, select scutter4.

10 Click Delete in the pop-up window if you are sure you want to remove the user account.

You are prompted to remove the user's home directory and mailbox contents.

Managing Client-Server Support (Overview)

This chapter describes the management of server and client support on a network. Overview information is provided about each system configuration (referred to as a *system type*) that is supported in the Oracle Solaris OS. This chapter also includes guidelines for selecting the appropriate system type to meet your needs.

Note – You cannot use the `smosservice` and the `smdiskless` commands on systems that have an Oracle Solaris ZFS root file system installed. This is a known issue with all Solaris releases that support the installation of a ZFS root file system.

You can quickly provision systems that run a UFS root file system or a ZFS root file system by using the Solaris Flash installation feature. For more information, see [“Installing a ZFS Root File System \(Oracle Solaris Flash Archive Installation\)”](#) in *Oracle Solaris ZFS Administration Guide*.

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

- “What's New in Managing Client-Server Support?” on page 132
- “Where to Find Client-Server Tasks” on page 133
- “What Are Servers, Clients, and Appliances?” on page 134
- “What Does Client Support Mean?” on page 134
- “Overview of System Types” on page 135
- “Diskless Client Management Overview” on page 138

For step-by-step instructions about how to manage diskless client support, see [Chapter 7](#), “Managing Diskless Clients (Tasks).”

What's New in Managing Client-Server Support?

This section describes new or changed diskless client features in this Solaris release. For a complete listing of new features and a description of Oracle Solaris releases, see *Oracle Solaris 10 8/11 What's New*.

Support for Specifying Platform by Using the `bootadm -p` Command

A new `-p platform` argument has been added to the `bootadm` command. This option enables you to specify the platform or machine hardware class of a client system in situations where the client platform differs from the server platform, for example when administering diskless clients.

For more information, see the `bootadm(1M)` man page.

`nfs4_domain` Keyword Impacts Diskless Client Boot

The `set_nfs4_domain` script that was delivered in Oracle Solaris 10 is no longer used to set the NFSv4 domain. To set the NFSv4 domain, add the new `nfs4_domain` keyword to the diskless client's `sysidcfg` file. Note that if the `nfs4_domain` keyword exists in the `sysidcfg` file, the first boot of a diskless client sets the domain accordingly.

x86: Diskless Client Changes That Apply to GRUB

The following feature enhancements are part of the new diskless boot scheme:

- The OS server is now capable of serving multiple Solaris releases simultaneously.
 - With the new diskless boot scheme, you can perform a `pxegrub` based network boot, where multiple releases are presented to a client from the GRUB menu.
- Vendor-specific options are now specified in the boot archive.

In previous releases, client-specific boot properties, typically defined in the `bootenv.rc` file, were provided by using vendor-specific options for the DHCP setup. The total length of the information that was required frequently exceeded the limit in the DHCP specification.

With the new boot scheme, this information is now part of the boot archive. The PXE/DHCP server is only required to provide the server IP address, the boot file, `pxegrub`, and possibly a client-specific menu file, through Site Option 150.

x86: Changes to the `smdiskless` Command

The `smdiskless` command is used to set up diskless clients. Previously, the `smdiskless` command set up the root (`/`) and `/usr` file systems, then exported these file systems to the client through NFS. To boot the client, you would additionally need to configure the `/tftpboot` area manually. This manual step is no longer a requirement for setting up a diskless client. The `smdiskless` command now automatically invokes a script in the `/usr/sadm/lib/wbem/config_tftp` file, which prepares the `/tftpboot` area for a PXE boot.

After running the `smdiskless` command, the `/tftpboot/01ethernet-address` file is displayed as a link to `pxegrub` and the `/tftpboot/menu.lst.01ethernet-address` file, which contains the GRUB menu entry. The `ethernet-address` in this instance is 01, followed by the Ethernet address of the client network interface. When supplying the Ethernet address of the client network interface, use uppercase letters and do not include colons.

The boot archive of the diskless client is automatically updated during shutdown. If the client's boot archive is out of date when it is shut down, you might need to run the following command from the OS server to update the boot archive:

```
# bootadm update-archive -f -R /export/root/host-name
```

where *host-name* is the host name of the client system.

For more information, see [“x86: How to Boot in Failsafe Mode to Forcibly Update a Corrupt Boot Archive” on page 258](#) and the `bootadm(1M)` man page.

Note – This information applies to both SPARC and x86 based OS servers that are serving x86 based clients.

For more information on setting up and configuring DHCP, see [Chapter 14, “Configuring the DHCP Service \(Tasks\)” in *System Administration Guide: IP Services*](#).

For more information on how to manage diskless clients in the GRUB boot environment, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245](#).

Where to Find Client-Server Tasks

Use this table to find step-by-step instructions for setting up server and client support.

Client-Server Services	For More Information
Install or JumpStart clients	Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations

Client-Server Services	For More Information
Diskless client systems in the Solaris OS	“Diskless Client Management Overview” on page 138 and Chapter 7, “Managing Diskless Clients (Tasks)”

What Are Servers, Clients, and Appliances?

Systems on the network can usually be described as one of the system types in this table.

System Type	Description
Server	A system that provides services to other systems in its network. There are file servers, boot servers, web servers, database servers, license servers, print servers, installation servers, appliance servers, and even servers for particular applications. This chapter uses the term <i>server</i> to mean a system that provides boot services and file systems for other systems on the network.
Client	<p>A system that uses remote services from a server. Some clients have limited disk storage capacity, or perhaps none at all. Such clients must rely on remote file systems from a server to function. Diskless systems and appliance systems are examples of this type of client.</p> <p>Other clients might use remote services (such as installation software) from a server. However, they don't rely on a server to function. A stand-alone system is a good example of this type of client. A stand-alone system has its own hard disk that contains the root (/), /usr, and /export/home file systems and swap space.</p>
Appliance	A network appliance such as the Sun Ray appliance provides access to applications and the Solaris OS. An appliance gives you centralized server administration, and no client administration or upgrades. Sun Ray appliances also provide <i>hot desking</i> . Hot desking enables you to instantly access your computing session from any appliance in the server group, exactly where you left off. For more information, see http://www.sun.com/software/index.jsp?cat=Desktop& .

What Does Client Support Mean?

Support can include the following:

- Making a system known to the network (host name and Ethernet address information)
- Providing installation services to remotely boot and install a system
- Providing Solaris OS services and application services to a system with limited disk space or no disk space

Overview of System Types

System types are sometimes defined by how they access the root (/) and /usr file systems, including the swap area. For example, stand-alone systems and server systems mount these file systems from a local disk. Other clients mount the file systems remotely, relying on servers to provide these services. This table lists some of the characteristics of each system type.

TABLE 6-1 Characteristics of System Types

System Type	Local File Systems	Local Swap Space?	Remote File Systems	Network Use	Relative Performance
Server	root (/) /usr /home /opt /export/home	Available	Not available	High	High
Stand-alone system	root (/) /usr /export/home	Available	Not available	Low	High
OS Server	/export/root				
Diskless client	Not available	Not available	root (/) swap /usr /home	High High	Low Low
Appliance	Not available	Not available	Not available	High	High

Description of a Server

A server system contains the following file systems:

- The root (/) and /usr file systems, plus swap space
- The /export and /export/home file systems, which support client systems and provide home directories for users
- The /opt directory or file system for storing application software

Servers can also contain the following software to support other systems:

- Oracle Solaris OS services for diskless systems that are running a different release



Caution – OS client-server configurations, where only one system is running a Solaris release that implements either GRUB on the x86 platform or the new boot architecture on the SPARC platform, can result in major incompatibilities. It is therefore recommended that you install or upgrade diskless systems to the same release as the server OS before adding diskless client support.

Note that GRUB based booting was introduced on the x86 platform in the Solaris 10 1/06 release. The new SPARC boot architecture was introduced in the Solaris 10 10/08 release.

- Clients that use a different platform than the server
- Oracle Solaris CD or DVD image software and boot software for networked systems to perform remote installations
- Oracle Solaris JumpStart directory for networked systems to perform custom JumpStart installations

Stand-Alone Systems

A *networked stand-alone system* can share information with other systems in the network. However, it can continue to function if detached from the network.

A stand-alone system can function autonomously because it has its own hard disk that contains the root (/), /usr, and /export/home file systems and swap space. Thus, the stand-alone system has local access to OS software, executables, virtual memory space, and user-created files.

Note – A stand-alone system requires sufficient disk space to hold its necessary file systems.

A *non-networked stand-alone system* is a stand-alone system with all the characteristics just listed, except it is not connected to a network.

Diskless Clients

A *diskless client* has no disk and depends on a server for all its software and storage needs. A diskless client remotely mounts its root (/), /usr, and /home file systems from a server.

A diskless client generates significant network traffic due to its continual need to procure OS software and virtual memory space from across the network. A diskless client cannot operate if it is detached from the network or if its server malfunctions.

For more overview information about diskless clients, see [“Diskless Client Management Overview”](#) on page 138.

Description of an Appliance

An appliance, such as the Sun Ray appliance, is an X display device that requires no administration. There is no CPU, fan, disk, and very little memory. An appliance is connected to a Sun display monitor. However, the appliance user's desktop session is run on a server and displayed back to the user.

The X environment is set up automatically for the user and has the following characteristics:

- Relies on a server to access other file systems and software applications
- Provides centralized software administration and resource sharing
- Contains no permanent data, making it a field-replaceable unit (FRU)

Guidelines for Choosing System Types

You can determine which system types are appropriate for your environment by comparing each system type based on the following characteristics:

Centralized administration:

- Can the system be treated as a field-replaceable unit (FRU)?
This means that a broken system can be quickly replaced with a new system without any lengthy backup and restore operations and no loss of system data.
- Does the system need to be backed up?
Large costs in terms of time and resources can be associated with backing up a large number of desktop systems.
- Can the system's data be modified from a central server?
- Can the system be installed quickly and easily from a centralized server without handling the client system's hardware?

Performance

- Does this configuration perform well in desktop usage?
- Does the addition of systems on a network affect the performance of other systems already on the network?

Disk space usage

How much disk space is required to effectively deploy this configuration?

This table describes how each system type scores in terms of each characteristic. A ranking of 1 is most efficient. A ranking of 4 is least efficient.

TABLE 6-2 Comparison of System Types

System Type	Centralized Administration	Performance	Disk Space Usage
Stand-alone system	4	1	4
Diskless client	1	4	1
Appliance	1	1	1

Diskless Client Management Overview

The following sections and [Chapter 7, “Managing Diskless Clients \(Tasks\)”](#), describe how to manage diskless client support in the Oracle Solaris OS.

A *diskless client* is a system that depends on an *OS server* for its operating system, software, and storage. A diskless client mounts its root (*/*), */usr*, and other file systems from its OS server. A diskless client has its own CPU and physical memory and can process data locally. However, a diskless client cannot operate if it is detached from its network or if its OS server malfunctions. A diskless client generates significant network traffic because of its continual need to function across the network.

Starting with the Solaris 9 release, the diskless client commands, `smosservice` and `smdiskless`, enable you to manage OS services and diskless client support. In the Solaris 8 release, diskless clients were managed with the Solstice GUI management tools.

OS Server and Diskless Client Support Information



Caution – Attempts to add diskless client support using an OS client-server configuration where one system implements the new boot architecture, but the other system does not, can result in major incompatibilities. New boot (GRUB) was implemented on the x86 platform, starting with the Solaris 10 1/06 release and on the SPARC platform, starting with the Solaris 10 10/8 release. Note that adding diskless support on systems that are running a Solaris release that is more recent than that which is running on the OS server is also an unsupported configuration. To avoid potential problems, it is recommended that you install the latest Solaris release before adding diskless client support.

The Solaris releases and architecture types that are supported by the `smosservice` and `smdiskless` commands include the following:

- **SPARC based servers:** Supported in the Solaris 8, Solaris 9, and Solaris 10 releases
- **SPARC based clients:** Supported in the Solaris 8, Solaris 9, and Solaris 10 releases
- **x86 based servers:** Supported in the Solaris 9, and Solaris 10 releases

- **x86 based clients:** Supported in the Solaris 9, and Solaris 10 releases

The following table shows the x86 OS client-server configurations that are supported by the `smosservice` and `smdiskless` commands. This information applies to the Solaris 9 and the Oracle Solaris 10 FCS (3/05) release.

If you are running at least the Solaris 10 1/06 release, it is recommended that you install or upgrade to the same release before adding diskless client support.

TABLE 6-3 x86 OS Client-Server Support

Server OS	Diskless Client OS	
	Oracle Solaris 10 3/05	Solaris 9
Oracle Solaris 10 3/05	Supported	Supported
Solaris 9	Not supported	Supported

The following table shows the SPARC OS client-server configurations that are supported by the `smosservice` and `smdiskless` commands. This information applies to the Solaris 8 and Solaris 9 releases, and the Oracle Solaris OS, up through the 10 5/08 OS.

If you are running at least the Solaris 10 10/08 release, it is recommended that you install or upgrade to the same release before adding diskless client support.

TABLE 6-4 SPARC OS Client-Server Support

Server OS	Diskless Client OS		
	Oracle Solaris 10 3/05 through Solaris 10 5/08	Solaris 9	Solaris 8
Oracle Solaris 10 3/05 through Solaris 10 5/08	Supported	Supported	Supported
Solaris 9	Not supported	Supported	Supported
Solaris 8	Not supported	Not supported	Supported

Diskless Client Management Features

You can use the `smosservice` and `smdiskless` commands to add and maintain diskless client support on a network. By using a name service, you can manage system information in a centralized manner so that important system information, such as host names, do not have to be duplicated for every system on the network.

You can perform the following tasks with the `smossservice` and `smdiskless` commands:

- Add and modify diskless client support
- Add and remove OS services
- Manage diskless client information in the LDAP, NIS, NIS+, or files name service environment

If you are performing a GRUB based boot on an x86 system, you need to manually set up the DHCP configuration. See [“x86: How to Prepare for Adding Diskless Clients in a GRUB Based Boot Environment”](#) on page 146 for more information.

Note – You can only use the diskless client commands to set up diskless client booting. You cannot use these commands to set up other services, such as remote installation or profile services. Set up remote installation services by including diskless client specifications in the `sysidcfg` file. For more information, see *Oracle Solaris 10 8/11 Installation Guide: Custom JumpStart and Advanced Installations*.

Working With Diskless Client Commands

By writing your own shell scripts and using the commands shown in the following table, you can easily set up and manage your diskless client environment.

TABLE 6-5 Diskless Client Commands

Command	Subcommand	Task
<code>/usr/sadm/bin/smossservice</code>	<code>add</code>	Add OS services
	<code>delete</code>	Delete OS services
	<code>list</code>	List OS services
	<code>patch</code>	Manage OS service patches
<code>/usr/sadm/bin/smdiskless</code>	<code>add</code>	Add a diskless client to an OS server
	<code>delete</code>	Delete a diskless client from an OS server
	<code>list</code>	List the diskless clients on an OS server
	<code>modify</code>	Modify the attributes of a diskless client

You can obtain help on these commands in two ways:

- Use the `-h` option when you type the command, subcommand, and required options, as shown in the following example.

```
% /usr/sadm/bin/smdiskless add -p my-password -u my-user-name -- -h
```

- View the [smdiskless\(1M\)](#) and [smosservice\(1M\)](#) man pages.

Required RBAC Rights for Diskless Client Management

You can use the `smosservice` and `smdiskless` commands as superuser. If you are using role-based access control (RBAC), you can use either a subset of or all of the diskless client commands, according to the RBAC rights to which they are assigned. The following table lists the RBAC rights that are required to use the diskless client commands.

TABLE 6-6 Required RBAC Rights for Diskless Client Management

RBAC Right	Command	Task
Basic Solaris User, Network Management	<code>smosservice list</code>	List OS services
	<code>smosservice patch</code>	List OS service patches
	<code>smdiskless list</code>	List diskless clients on an OS server
Network Management	<code>smdiskless add</code>	Add diskless clients
System Administrator	All commands	All tasks

Adding OS Services

An Oracle Solaris OS server is a server that provides operating system (OS) services to support diskless client systems. You can add support for an OS server or convert a stand-alone system to an OS server by using the `smosservice` command.

For each platform group and Oracle Solaris release that you want to support, you must add the particular OS service to the OS server. For example, if you want to support SPARC sun-4u systems running Oracle Solaris , you must add sun-4u/Oracle Solaris 10 OS services to the OS server. For each diskless client that you support, you must add the OS service for that client to the OS server. For example, you would need to add OS services to support SPARC sun-4m systems or x86 based systems that run Oracle Solaris 10 or the Solaris 9 release, because they are different platform groups.

You must have access to the appropriate Oracle Solaris software CD, DVD, or disk image to add OS services.

Adding OS Services When the OS Server Has Been Patched

When adding OS services to an OS server, you might see an error message stating that you have inconsistent versions of the OS running on the server and the OS that you are trying to add. This error message occurs when the installed version of the OS has packages that were previously patched, and the OS services being added do not have those packages patched, because the patches have been integrated into the packages.

For example, you might have a server that is running the current Solaris release or the Oracle Solaris OS. You might also have additional OS services loaded on this server, including the Solaris 9 SPARC sun-4m OS services that have been patched. If you try to add the Solaris 8 SPARC sun-4u OS services from a CD-ROM to this server, you could get the following error message:

```
Error: inconsistent revision, installed package appears to have been
patched resulting in it being different than the package on your media.
You will need to backout all patches that patch this package before
retrying the add OS service option.
```

Disk Space Requirements for OS Servers

Before you set up your diskless client environment, ensure that you have the required disk space available for each diskless client directory.

In previous Solaris releases, you were prompted about diskless client support during the installation process. Starting with the Solaris 9 release, you must manually allocate an `/export` file system either during installation or create it after installation. See the following table for specific disk space requirements.

TABLE 6-7 Disk Space Recommendations for Solaris OS Servers and Diskless Clients

Server OS/Architecture Type	Directory	Required Disk Space
Oracle Solaris 10 SPARC based OS server	<code>/export</code>	5 to 6.8 Gbytes
Oracle Solaris 10 x86 based OS server	<code>/export</code>	5 to 6.8 Gbytes
Oracle Solaris 10 SPARC based diskless client	<code>/export</code>	Reserve 200 to 300 Mbytes per diskless client.
Oracle Solaris 10 x86 based diskless client	<code>/export</code>	Reserve 200 to 300 Mbytes per diskless client.

Note – Disk space recommendations can vary, depending on the Oracle Solaris release that is installed. For specific information about the disk space recommendations in the current Solaris release, see “Disk Space Recommendations for Software Groups” in *Oracle Solaris 10 8/11 Installation Guide: Planning for Installation and Upgrade*.

Managing Diskless Clients (Tasks)

This chapter describes how to manage diskless clients in the Oracle Solaris OS.

Note – You cannot use the `smosservice` and the `smdiskless` commands on systems that have an Oracle Solaris ZFS root file system installed. This is a known issue with all Solaris releases that support the installation of a ZFS root file system.

You can quickly provision systems that run a UFS root file system or a ZFS root file system by using the Solaris Flash installation feature. For more information, see [“Installing a ZFS Root File System \(Oracle Solaris Flash Archive Installation\)”](#) in *Oracle Solaris ZFS Administration Guide*.

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

- “Managing Diskless Clients (Task Map)” on page 143
- “Preparing for Managing Diskless Clients” on page 145
- “Patching Diskless Client OS Services” on page 160
- “Troubleshooting Diskless Client Installation Problems” on page 163

For overview information on managing diskless clients, see [Chapter 6, “Managing Client-Server Support \(Overview\)”](#).

Managing Diskless Clients (Task Map)

The following table identifies the procedures that are required to manage diskless clients.

Task	Description	For Instructions
1. (Optional) Enable Solaris Management Console logging to view diskless client error messages.	Choose Log Viewer from the console main window to view diskless client error messages.	“Starting the Solaris Management Console” on page 44

Task	Description	For Instructions
2. Prepare for adding a diskless client.	Verify supported releases and identify the platform, media path, and cluster (or software group) of each diskless client.	<p>“x86: How to Prepare for Adding Diskless Clients in a GRUB Based Boot Environment” on page 146</p> <p>“How to Prepare for Adding Diskless Clients in Oracle Solaris 10” on page 148</p>
3. Add required OS services to an OS server.	Add the OS services for the diskless clients you want to support by using the <code>smosservice</code> command. You must identify the platform, media path, and each diskless client platform that you want to support.	“How to Add OS Services for Diskless Client Support” on page 150
4. Locate and install any ARCH=all packages that were missed when you added OS services to the server. Note – To avoid having to add these packages to each diskless client individually, perform this task <i>prior</i> to adding diskless client support.	The <code>smosservice add</code> command does not install any root (/) or <code>/usr</code> packages that are designated ARCH=all. These packages must be installed manually after adding the OS services to the OS server.	“How to Locate and Install Missing ARCH=all Packages” on page 163
5. Add a diskless client.	Add diskless client support by specifying all required information with the <code>smdiskless</code> command.	<p>“x86: How to Add a Diskless Client in the GRUB Based Boot Environment” on page 152</p> <p>“How to Add a Diskless Client in Oracle Solaris 10” on page 155</p>
6. Boot the diskless client.	Verify that a diskless client was successfully added by booting the diskless client.	<p>“x86: How to Boot a Diskless Client With GRUB” on page 157</p> <p>“SPARC: How to Boot a Diskless Client in Oracle Solaris 10” on page 158</p>
7. (Optional) Delete diskless client support.	Delete support for a diskless client, if it is no longer required.	“How to Remove Diskless Client Support” on page 158
8. (Optional) Delete OS services for a diskless client.	Delete OS services for a diskless client, if they are no longer needed.	“How to Remove OS Services for Diskless Clients” on page 159
9. (Optional) Patch OS services.	Add, delete, list, or synchronize patches for diskless client OS services.	“How to Add an OS Patch for a Diskless Client” on page 161

Preparing for Managing Diskless Clients

The following sections describe the preparations that are necessary for managing diskless clients.

Keep the following key points in mind when managing diskless clients:

- The Oracle Solaris installation program doesn't prompt you to set up diskless client support. You must manually create an `/export` partition to support diskless clients. You create the `/export` partition during or after the installation process.
- The `/export` partition must contain a minimum of 5 Gbytes, depending upon the number of clients supported. For specific information, see *Disk Space Requirements for OS Servers*.
- The name service identified in the `smossservice` or `smdiskless` commands must match the primary name service identified in the `/etc/nsswitch.conf` file. If you don't specify a name service in the `smdiskless` or `smossservice` commands, the default name service is `files`.

Use the `-D` option to the `smossservice` and `smdiskless` commands to specify a name server. For more information, see the [smossservice\(1M\)](#) and [smdiskless\(1M\)](#) man pages.

Starting with the Solaris 10 8/07 release, the `set_nfs4_domain` script that was delivered in Oracle Solaris 10 is no longer used to set the NFSv4 domain. To set the NFSv4 domain, add the `nfs4_domain` keyword to the diskless client's `sysidcfg` file, for example, `server:/export/root/client/etc/sysidcfg`.

If the `nfs4_domain` keyword exists in the client system's `sysidcfg` file, the first boot of a diskless client sets the domain accordingly. Also, the OS server should be up and running, and the diskless client's NFSv4 domain setting must match the setting in the OS server's `/var/run/nfs4_domain` file.

For more information, see “[Preconfiguring With the sysidcfg File](#)” in *Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations*.

- The OS server and the diskless client must be on the same subnet.
- You cannot provide client services on a multiterabyte UFS file system, because OS and diskless client services *cannot* be added to a UFS file system that resides on an EFI-labeled disk.

Note – Attempts to add OS and diskless client services to a UFS file system that resides on an EFI-labeled disk result in an erroneous insufficient disk space message similar to the following:

```
The partition /export does not have enough free space.
1897816 KB (1853.34 MB) additional free space required.
Insufficient space available on
/dev/dsk/c0t5d0s0 /export
```

After you determine the platform, media path, and cluster for each diskless client, you are ready to add OS services.

The following directories are created and populated for each OS service that you add:

- `/export/Solaris_version/Solaris_version-instruction-set.all` (symbolic link to `/export/exec/Solaris_version/Solaris_version-instruction-set.all`)
- `/export/Solaris_version`
- `/export/Solaris_version/var`
- `/export/Solaris_version/opt`
- `/export/share`
- `/export/root/templates/Solaris_version`
- `/export/root/clone`
- `/export/root/clone/Solaris_version`
- `/export/root/clone/Solaris_version/ machine-class`

The following default directories are created and populated on the OS server for each diskless client that you add:

- `/export/root/diskless-client`
- `/export/swap/diskless-client`
- `/tftpboot/diskless-client-ipaddress-in-hex/export/dump/diskless-client` (if you specify the `-x dump` option)

Note – You can modify the default locations of the `/`, `/swap`, and `/dump` directories by using the `-x` option to the `smoservice` and `smdiskless` commands. However, do not create these directories under the `/export` file system.

▼ **x86: How to Prepare for Adding Diskless Clients in a GRUB Based Boot Environment**

Use this procedure to prepare for adding a diskless client. This procedure includes general information for x86 based systems.

When you use the `smoservice add` command to add OS services, you must specify the platform, media path, and cluster (or software group) of each diskless client platform that you want to support.

Before You Begin Ensure that the system that is intended to be the OS service is running a supported release. Also, verify that the OS server release and diskless client release combination is supported. For more information, see “[OS Server and Diskless Client Support Information](#)” on page 138.

1 Identify the diskless client platform by using this format:

instruction-set.machine-class.Solaris-version

For example:

`i386.i86pc.Solaris_10`

The following are the possible platform options:

Instruction Set	Machine Class	Solaris Version
sparc	sun4v	Starting with the Solaris 10 1/06 OS
	sun4u, sun4m, sun4d, and sun4c	Oracle Solaris 10, Solaris 9, and Solaris 8
	i386	Oracle Solaris 10, Solaris 9, and Solaris 8

Note – The sun-4c architecture is not supported in the Solaris 8, Solaris 9, or Oracle Solaris 10. The sun-4d architecture is not supported in the Solaris 9 or Oracle Solaris 10. The sun-4m architecture is not supported in the Oracle Solaris 10 OS.

2 Identify the media path.

The media path is the full path to the disk image that contains the OS that you want to install for the diskless client.

In some of the Oracle Solaris releases, the OS is delivered on multiple CDs. However, you cannot use the `smosservice` command to load OS services from a multiple CD distribution. You must run the scripts that are found on the Oracle Solaris software CDs (and optional Language CD), or the Oracle Solaris DVD, as described in the steps that follow:

Note – In this Oracle Solaris release, the software is delivered on DVD *only*.

3 Create an install image on a server.

For information on setting up an installation server, refer to [Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations](#).

4 Load the required OS services from the DVD image.

```
# /mount_point/Solaris_10/Tools/setup_install_server
```

5 Add the `BootFile` and `BootSrvA` DHCP options to your DHCP server configuration to enable a PXE boot.

For example:

```
Boot server IP (BootSrvA) : svr-addr  
(BootFile) : 01client-macro
```

where *svr-addr* is the IP address of the OS server and *client-macro* is named by the client's Ethernet type (01) and the Media Access Control (MAC) address of the client. This number is also the name of the file that is used in the `/tftpboot` directory on the installation server.

Note – The notation for the *client-macro* consists of uppercase letters. This notation should not contain any colons.

You can add these options from the command-line, or by using DHCP Manager. See [Example 7–4](#) for more information.

For more information, see “x86: How to Perform a GRUB Based Boot From the Network” on page 263, “Preconfiguring System Configuration Information With the DHCP Service (Tasks)” in *Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations*, and Part III, “DHCP” in *System Administration Guide: IP Services*.

6 After the Oracle Solaris image is installed on the disk, note the disk media path. For example:

```
/net/export/install/sol_10_x86
```

This is the disk media path that needs to be specified when you use the `smos service` command.

7 Identify the `SUNWCxall` cluster when you add OS services.

You must use the same cluster for diskless clients that run the same OS on the same system.

Note – Always specify `SUNWCxall` as the cluster.

▼ How to Prepare for Adding Diskless Clients in Oracle Solaris 10

When you use the `smos service add` command to add OS services, you must specify the platform, media path, and cluster (or software group) of each diskless client platform that you want to support.

Before You Begin Ensure that the system that is intended to be the OS service is running a supported release. Also verify that the combination of OS server release and diskless client release is supported. For more information, see “OS Server and Diskless Client Support Information” on page 138.

1 Identify the diskless client platform by using this format:

instruction-set.machine-class.Solaris- version

For example:

`sparc.sun4u.Solaris_10`

The following are the possible platform options:

<i>instruction-set</i>	<i>machine-class</i>	<i>Solaris_version</i>
sparc	sun4v	Starting with the Solaris 10 1/06 OS
	sun4c, sun4d, sun4m, sun4u,	Solaris_10, Solaris_9, and Solaris_8
i386	i86pc	Solaris_10, Solaris_9, and Solaris_8

Note – The sun-4c architecture is not supported in the Solaris 8, Solaris 9, or Oracle Solaris 10. The sun-4d architecture is not supported in the Solaris 9 or Oracle Solaris 10. The sun-4m architecture is not supported in the Oracle Solaris OS.

2 Identify the media path.

The media path is the full path to the disk image that contains the OS that you want to install for the diskless client.

In some Oracle Solaris releases, the OS is delivered on multiple CDs. However, you cannot use the `smosservice` command to load OS services from a multiple CD distribution. You must run the scripts that are found on the Solaris software CDs (and optional Language CD), or the DVD, as described in the steps that follow:

Note – In this Oracle Solaris release, the software is delivered on DVD *only*.

3 Create an install image on a server.

For information on setting up an installation server, refer to [Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations](#).

4 Load the required OS services from the DVD image.

```
# /mount_point/Solaris_10/Tools/setup_install_server
```

5 After the Oracle Solaris DVD image is installed on the disk, specify the disk media path. For example:

```
/export/install/sparc_10
```

6 Identify the SUNWCxall cluster when you add OS services.

You must use the same cluster for diskless clients that run the same OS on the same system. For example, consider the following Solaris 9 diskless clients:

- `sparc.sun4m.Solaris_9`
- `sparc.sun4u.Solaris_9`

To set up these diskless clients, you would need to specify the SUNWCxall cluster for each diskless client because the sun4u and sun4m systems require the SUNWCxall cluster. In addition, diskless clients that run the same operating release (in this example, Solaris_9) on the same system must use the same cluster.

Note – If you are using a sun4u system, or if you are using a system with an accelerated 8-bit color memory frame buffer (`cgsix`), you *must* specify SUNWCxall as the cluster.

▼ How to Add OS Services for Diskless Client Support

Use this procedure to add OS services for a diskless client on the server.

Note – When adding OS services with the `smos service add` command, root (/) and /usr packages with the ARCH=all type are not installed. These packages are skipped. No warning or error messages are displayed. After you add the OS services to the OS server, you must install the missing packages manually. For instructions, see [How to Locate and Install Missing ARCH=all Packages](#).

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Verify that the Solaris Management Console server is running and that the diskless client tools are available on the system.

```
# /usr/sadm/bin/smosservice list -H host-name:898 --
```

3 Add the OS services.

```
# /usr/sadm/bin/smosservice add -H host-name
:898 -- -o
host-name
-x mediapath=path -x platform=
instruction-set.machine-class
.Solaris_version
-x cluster=cluster-name
-x locale=locale-name
add
```

Adds the specified OS service.

`-H host-name:898`

Specifies the host name and port to which you want to connect. If you do not specify a port, the system connects to the default port, 898.

Note – The `-H` option is not a required option when using the `smossservice` command to add OS services.

--

Identifies that the subcommand arguments start after this point.

`-x mediapath=path`

Specifies the full path to the Solaris image.

`-x platform=instruction-set.machine-class.Solaris_version`

Specifies the instruction architecture, machine class, and the Solaris version to be added.

`-x cluster=cluster-name`

Specifies the Solaris cluster to install.

`-x locale=locale-name`

Specifies the locale to install.

Note – The installation process can take about 45 minutes, depending on the server speed and the OS service configuration you choose.

For more information, see the [smossservice\(1M\)](#) man page.

- 4 (Optional) Continue to add the other OS services.
- 5 When you are finished adding OS services, verify that the OS services were installed.

```
# /usr/sadm/bin/smossservice list -H host-name:898 --
```

Example 7-1 SPARC: Adding an OS Service for Diskless Client Support

This example shows how to add Solaris 10 SPARC based OS services on the server `jupiter`. The server `jupiter` is running the Oracle Solaris OS. The CD image of the Oracle Solaris 10 SPARC based OS is located on the installation server, `myway`, in `/export/s10/combined.s10s_u2wos/61`.

```
# /usr/sadm/bin/smossservice add -H jupiter:898 -- -o jupiter
-x mediapath=/net/myway/export/s10/combined.s10s_u2wos/61
-x platform=sparc.sun4u.Solaris_10
-x cluster=SUNWCXall -x locale=en_US
```

```
# /usr/sadm/bin/smossservice list - H jupiter:898
Authenticating as user: root
```

```
Type /? for help, pressing enter accepts the default denoted by [ ]
Please enter a string value for: password :: xxxxxx
Loading Tool: com.sun.admin.osservmgr.cli.OsServerMgrCli
from jupiter:898
Login to jupiter as user root was successful.
Download of com.sun.admin.osservmgr.cli.OsServerMgrCli from jupiter:898
was successful.
```

Example 7-2 x86: Adding an OS Service for Diskless Client Support

This example shows how to add Solaris 10 x86 based OS services on the server orbit. The server orbit is running the Oracle Solaris OS. The CD image of the Oracle Solaris 10 x86 based OS is located on the installation server, seriously, in `/export/s10/combined.s10x_u2wos/03`.

```
# /usr/sadm/bin/smosservice add -H orbit:898 -- -o orbit -x
mediapath=/net/seriously/export/s10u2/combined.s10x_u2wos/03 -x
platform=i386.i86pc.Solaris_10 -x cluster=SUNWCXall -x locale=en_US

# /usr/sadm/bin/smosservice list - H orbit:898
Type /? for help, pressing <enter> accepts the default denoted by [ ]
Please enter a string value for: password ::
Starting Solaris Management Console server version 2.1.0.
endpoint created: :898
Solaris Management Console server is ready.
Loading Tool: com.sun.admin.osservmgr.cli.OsServerMgrCli from orbit:898
Login to orbit as user root was successful.
Download of com.sun.admin.osservmgr.cli.OsServerMgrCli from orbit:898 was successful.
Client          Root Area
                Swap Area
                Dump Area
-----
.
.
.
#
```

Next Steps Locate and install any ARCH=all packages that were missed when you ran the `smosservice add` command to add the OS services to the OS server. For more information, see [How to Locate and Install Missing ARCH=all Packages](#).

▼ x86: How to Add a Diskless Client in the GRUB Based Boot Environment

Starting with the Solaris 10 1/06 release, use this procedure to add a diskless client after you have added OS services.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Add the diskless client.

```
# /usr/sadm/bin/smdiskless add -- -i
ip-address -e ethernet-address
-n client-name -x os=
instruction-set.machine-class.Solaris_
version
-x root=/export/root/client-name
-x swap=/export/swap/client-name
-x swapsize=size -x tz=
time-zone -x locale=
locale-name
```

add

Adds the specified diskless client.

--

Identifies that the subcommand arguments start after this point.

-i *ip-address*

Identifies the IP address of the diskless client.

-e *ethernet-address*

Identifies the Ethernet address of the diskless client.

-n *client-name*

Specifies the name of the diskless client.

-x *os=instruction-set.machine-class.Solaris_version*

Specifies the instruction architecture, machine class, OS, and the Solaris version for the diskless client.

-x *root=root=/export/root/ client-name*

Identifies the root (/) directory for the diskless client.

-x *swap=root=/export/root/ client-name*

Identifies the swap file for the diskless client.

-x *swapsize=size*

Specifies the size of the swap file in Mbytes. The default is 24 Mbytes.

-x *tz=time-zone*

Specifies the time-zone for the diskless client.

-x *locale=locale-name*

Specifies the locale to install for the diskless client.

For more information, see the `smdiskless(1M)` man page.

3 If not already created, add the `BootSrvA` and `BootFile` DHCP options to your DHCP server configuration to enable a PXE boot.

For example:

```
Boot server IP (BootSrvA) : svr-addr
Boot file (BootFile) : 01client-macro
```

where *svr-addr* is the IP address of the server and *client-macro* is named by the client's Ethernet type (01) and the MAC address of the client. This number is also the name of the file that is used in the `/tftpboot` directory on the installation server.

Note – The *client-macro* notation consists of uppercase letters. The notation should not contain any colons.

The following files and directories are created in the `/tftpboot` directory:

```
drwxr-xr-x  6 root sys      512 Dec 28 14:53 client-host-name
lrwxrwxrwx  1 root root      31 Dec 28 14:53 menu.lst.01ethernet-address
                -> /tftpboot/client-host-name/grub/menu.lst
-rw-r--r--  1 root root 118672 Dec 28 14:53 01ethernet-address
```

- 4 If the console is on a serial port, edit the `/tftpboot/menu.lst.01ethernet-address` file to uncomment the line that specifies the `tty` setting.
 - To change the default `menu.lst` file that is created on the client, edit the `echo` lines in the `/usr/sadm/lib/wbem/config_tftp` file.

For more information, see “[Booting an x86 Based System From the Network](#)” on page 260.

- 5 Verify that the diskless clients were installed.


```
# /usr/sadm/bin/smdiskless list -H host-name:898 --
```
- 6 (Optional) Continue to use the `smdiskless add` command to add each diskless client.

Example 7–3 x86: Adding Diskless Client Support to an x86 Based System in the GRUB Boot Environment

This example shows how to add a Solaris 10 x86 based diskless client, `mikey1`.

```
rainy-01# /usr/sadm/bin/smdiskless add -H sdts-01-qfe0 -- -o sdts-01-qfe0
-n mikey1 -i 192.168.20.22 -e 00:E0:88:55:33:BC -x os=i386.i86pc.Solaris_10
-x root=/export/root/mikey1 -x swap=/export/swap/mikey1
```

```
Loading Tool: com.sun.admin.ossvermgr.cli.OsServerMgrCli
from sdts-01-qfe0
Login to rainy-01-qfe0 as user root was successful.
Download of com.sun.admin.ossvermgr.cli.OsServerMgrCli from
rainy-01-qfe0 was successful.
```

```
# /usr/sadm/bin/smdiskless list -H mikey1:898 --
Loading Tool: com.sun.admin.ossvermgr.cli.OsServerMgrCli from mikey1:898
Login to mikey1 as user root was successful.
Download of com.sun.admin.ossvermgr.cli.OsServerMgrCli from mikey1:898 was
```

```

successful.
Platform
-----
i386.i86pc.Solaris_10
sparc.sun4us.Solaris_10
sparc.sun4u.Solaris_10
i386.i86pc.Solaris_9

```

Example 7-4 x86: Adding the BootSrvA and BootFile DHCP Options to the DHCP Server Configuration

This example shows how to add the BootSrvA and BootFile DHCP options that are necessary for enabling a PXE boot.

```

rainy-01# pntadm -A mikey1 -m 0100E0885533BC -f 'MANUAL+PERMANENT' \
-i 0100E0885533BC 192.168.0.101

```

```

rainy-01# dhtadm -A -m 0100E0885533BC -d \
":BootSrvA=192.168.0.1:BootFile=0100E0885533BC:"

```

In the preceding examples, the server address is the IP address of the server, and the client macro is named by the client's Ethernet type (01) and its MAC address. This number is also the name of the file that is used in the `/tftpboot` directory on the installation server. Note that the notation for the client macro consists of uppercase letters and should not contain any colons.

▼ How to Add a Diskless Client in Oracle Solaris 10

Use this procedure to add a diskless client after you have added OS services. Unless otherwise noted, this procedure includes general information for both the SPARC and x86 platforms.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Add the diskless client.

```

# /usr/sadm/bin/smdiskless add -- -i
ip-address -e ethernet-address
-n client-name -x os=
instruction-set.machine-class.Solaris_
version
-x root=/export/root/client-name
-x swap=/export/swap/client-name
-x swappsize=size -x tz=
time-zone -x locale=
locale-name

```

add

Adds the specified diskless client.

- Identifies that the subcommand arguments start after this point.
- i *ip-address*
Identifies the IP address of the diskless client.
- e *ethernet-address*
Identifies the Ethernet address of the diskless client.
- n *client-name*
Specifies the name of the diskless client.
- x *os=instruction-set.machine-class .Solaris_version*
Specifies the instruction architecture, machine class, OS, and the Solaris version for the diskless client.
- x *root=root=/export/root/ client-name*
Identifies the root (/) directory for the diskless client.
- x *swap=root=/export/root/ client-name*
Identifies the swap file for the diskless client.
- x *swapsize=size*
Specifies the size of the swap file in Mbytes. The default is 24 Mbytes.
- x *tz=time-zone*
Specifies the time-zone for the diskless client.
- x *locale=locale-name*
Specifies the locale to install for the diskless client.

For more information, see the [smdiskless\(1M\)](#) man page.

- 3 (Optional) Continue to use the `smdiskless add` command to add each diskless client.
- 4 Verify that the diskless clients were installed.

```
# /usr/sadm/bin/smdiskless list -H host-name:898 --
```

Example 7-5 SPARC: Adding Diskless Client Support to a SPARC Based System

This example shows how to add Solaris 10 sun4u diskless client, `starlite`, from the server `bearclaus`.

```
# /usr/sadm/bin/smdiskless add -- -i 172.20.27.28 -e 8:0:20:a6:d4:5b
-n starlite -x os=sparc.sun4u.Solaris_10 -x root=/export/root/starlite
-x swap=/export/swap/starlite -x swapsize=128 -x tz=US/Mountain
-x locale=en_US

# /usr/sadm/bin/smdiskless list -H starlite:898 --
Loading Tool: com.sun.admin.ossvermgr.cli.OsServerMgrCli from line2-v480:898
Login to line2-v480 as user root was successful.
```

Download of com.sun.admin.ossvermgr.cli.OsServerMgrCli from line2-v480:898 was successful.
Platform

```
-----
i386.i86pc.Solaris_10
sparc.sun4us.Solaris_10
sparc.sun4u.Solaris_10
i386.i86pc.Solaris_9
sparc.sun4m.Solaris_9
sparc.sun4u.Solaris_9
sparc.sun4us.Solaris_9
```

Note that the `smdiskless list -H` command output lists both SPARC based and x86 based systems.

Example 7-6 x86: Adding Diskless Client Support to an x86 Based System That is Running Oracle Solaris 10

This example shows how to add an Oracle Solaris 10 x86 based diskless client, `mars`, from the server `bearclaus`.

```
# /usr/sadm/bin/smdiskless add -- -i 172.20.27.176 -e 00:07:E9:23:56:48
-n mars -x os=i386.i86pc.Solaris_10 -x root=/export/root/mars
-x swap=/export/swap/mars -x swapsize=128 -x tz=US/Mountain
-x locale=en_US
```

▼ x86: How to Boot a Diskless Client With GRUB

If you have installed or upgraded your system to at least the Solaris 10 1/06 OS, the procedure for booting a diskless client has changed. Follow these steps to boot a diskless client with GRUB.

Note – Starting with the Solaris 10 6/06 release, when booting the failsafe archive, you are no longer prompted by the system to automatically update the boot archives. The system prompts you to update the boot archives only if inconsistent boot archives are detected. For more information, see [“How to Boot an x86 Based System in Failsafe Mode”](#) on page 256.

Before You Begin To ensure that the system boots from the network, verify the following prerequisites on the OS server:

- Confirm that the name service used to add the diskless client and the OS services matches the primary name in the server's `/etc/nsswitch.conf` file.
- Verify that the DHCP and `tftp` boot services are running.
- Configure the system BIOS to boot from the network by enabling the PXE ROM option.
Some PXE-capable network adapters have a feature that enables PXE boot if you type a particular keystroke in response to a brief boot-time prompt. See your hardware documentation for information about how to set the boot priority in the BIOS.

1 Boot the diskless client by typing the correct keystroke combination.

The GRUB menu is displayed.

Depending on the configuration of your network installation server, the GRUB menu that is displayed on your system might vary from the GRUB menu that is shown here.

2 Use the arrow keys to select a boot entry, then press Enter.

If you do not make a selection, the default OS instance is automatically booted after several seconds.

- **If you need to modify the GRUB kernel behavior by editing the GRUB menu at boot time, use the arrow keys to select a boot entry, then type e to edit the entry.**

The boot command that you want to edit is displayed in the GRUB edit screen.

For more information about modifying kernel behavior at boot time, see [Chapter 11, “Modifying Oracle Solaris Boot Behavior \(Tasks\)”](#).

- **To save the edits and return to the previous menu, press Enter.**

The GRUB menu is displayed, showing the edits you made to the boot command.

- **Type b to boot the system from the network.**

▼ **SPARC: How to Boot a Diskless Client in Oracle Solaris 10**

Before You Begin Verify the following prerequisites on the OS server:

- Confirm that the name service used to add the diskless client and the OS services matches the primary name in the server's `/etc/nsswitch.conf` file.
Otherwise, the diskless client will not boot.
- Confirm that the `rpc.bootparamd` daemon is running. If it is not running, start it.

- **Boot the diskless client.**

```
ok boot net
```

▼ **How to Remove Diskless Client Support**

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

2 Remove the diskless client support.

```
# /usr/sadm/bin/smdiskless delete -- -o host-name
:898 -n client-name
```

3 Verify that the diskless client support has been removed.

```
# /usr/sadm/bin/smosservice list -H host-name:898 --
```

Example 7-7 Removing Diskless Client Support

This example shows how to remove the diskless client holoship from the OS server starlite.

```
# /usr/sadm/bin/smdiskless delete -- -o starlite:898 -n holoship
```

Authenticating as user: root

```
Type /? for help, pressing enter accepts the default denoted by [ ]
Please enter a string value for: password ::
Starting SMC server version 2.0.0.
endpoint created: :898
SMC server is ready.
```

```
# /usr/sadm/bin/smosservice list -H starlite:898 --
Loading Tool: com.sun.admin.osservermgr.cli.OsServerMgrCli from starlite
Login to starlite as user root was successful.
Download of com.sun.admin.osservermgr.cli.OsServerMgrCli from starlite
was successful.
```

▼ How to Remove OS Services for Diskless Clients**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Remove the OS services for the diskless clients.

```
# /usr/sadm/bin/smosservice delete -H $HOST:$PORT -u root -p $PASSWD --
-x instruction-set.all.Solaris_version
```

Note – Only the all machine-class is supported.

3 Verify that the OS services have been removed.

```
# /usr/sadm/bin/smosservice list -H host-name:898 --
```

Example 7-8 Removing OS Services for Diskless Clients

The following example shows how to removing the diskless client OS services (sparc.all.Solaris_10) from the server starlite.

```
# /usr/sadm/bin/smosservice delete -H starlite:898 -u root \  
-p xxxxxx -- -x sparc.all.solaris_10  
Authenticating as user: root  
Type /? for help, pressing enter accepts the default denoted by [ ]  
Please enter a string value for: password ::  
  
# /usr/sadm/bin/smosservice list -H starlite:898 --  
Loading Tool: com.sun.admin.osservmgr.cli.OsServerMgrCli from starlite:898  
Login to starlite as user root was successful.  
Download of com.sun.admin.osservmgr.cli.OsServerMgrCli from starlite:898  
was successful
```

Patching Diskless Client OS Services

Use the smosservice patch command to do the following:

- Establish the /export/diskless/Patches patch spool directory on an OS server.
- Add patches to the patch spool directory. If the patch you are adding obsoletes an existing patch in the spool, the obsolete patch is moved to /export/diskless/Patches/Archive.
- Delete patches from the patch spool directory.
- List the patches in the patch spool directory.
- Synchronize spooled patches out to clients. Note that you must reboot each synchronized client for the client to recognize the patch update.

Note – Keep your OS servers up to date by installing recommended OS patches on a timely basis.

For information on downloading patches, see [“Downloading a Patch” on page 432](#).

Displaying OS Patches for Diskless Clients

Diskless client patches are logged in different directories, depending on the type of patch:

- Kernel patches are logged in the diskless client's /var/sadm/patch directory. To display kernel patches, type the following command on the diskless client:

```
% patchadd -p
```

Note – You must be logged in to the diskless client when you run this command. Running the patchadd -p command on the OS server displays kernel patches for the OS server only.

- /usr patches are logged in the OS server's /export/Solaris_version/var/patch directory. A directory is created for each patch ID. To display /usr patches, type the following command on the OS server:


```
% patchadd -S Solaris_version -p
Patch: 111879-01 Obsolete: Requires: Incompatibles: Packages: SUNWwsr
```

To list all spooled patches by OS and architecture, use the `smosservice` command with the `-P` option.

▼ How to Add an OS Patch for a Diskless Client

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

2 Log in to the diskless client system and shut it down.

```
# init 0
```

3 Add the patch to a spool directory.

```
# /usr/sadm/bin/smosservice patch -- -a /var/patches/ patch-ID-revision
```

If the patch to add depends on another patch, adding the patch fails with the following message:

```
The patch patch-ID-revision could not be added
because it is dependent on other patches which have not yet been spooled.
You must add all required patches to the spool first.
```

4 Verify that the patch has been spooled.

```
# /usr/sadm/bin/smosservice patch -- -P
```

5 Push the spooled patch to the diskless client.

```
# /usr/sadm/bin/smosservice patch -- -m -U
```

Note – Pushing and synchronizing the patch to the diskless client can take up to 90 minutes per patch.

6 Verify the patch is applied to the diskless client.

```
# /usr/sadm/bin/smosservice patch -- -P
```

Example 7–9 Adding an OS Patch for a Diskless Client

This example shows how to add a Solaris 8 patch (111879-01) to the diskless client's OS services on the server.

```
# /usr/sadm/bin/smosservice patch -- -a /var/patches/111879-01
```

```
Authenticating as user: root
```

```

Type /? for help, pressing <enter> accepts the default denoted by [ ]
Please enter a string value for: password ::
Loading Tool: com.sun.admin.osservicepatch.cli.OsServerMgrCli from starlite
Login to starlite as user root was successful.
Download of com.sun.admin.osservicepatch.cli.OsServerMgrCli from starlite
was successful..
.
# /usr/sadm/bin/osservicepatch -- -P
Patches In Spool Area
Os Rel Arch Patch Id Synopsis
-----
8      sparc 111879-01 SunOS 5.8: Solaris Product Registry patch SUNWwsr

Patches Applied To OS Services
Os Service Patch
-----
Solaris_8

Patches Applied To Clone Areas
Clone Area Patch
-----
Solaris_8/sun4u Patches In Spool Area
Os Rel Arch Patch Id Synopsis
-----
8      sparc 111879-01 SunOS 5.8: Solaris Product Registry patch SUNWwsr
.
.
.
# /usr/sadm/bin/osservicepatch -- -m -U
Authenticating as user: root

Type /? for help, pressing <enter> accepts the default denoted by [ ]
Please enter a string value for: password ::
Loading Tool: com.sun.admin.osservicepatch.cli.OsServerMgrCli from starlite
Login to starlite as user root was successful.
Download of com.sun.admin.osservicepatch.cli.OsServerMgrCli from starlite
was successful.

# /usr/sadm/bin/osservicepatch -- -P
Authenticating as user: root
.
.
.
Patches In Spool Area
Os Rel Arch Patch Id Synopsis
-----
8      sparc 111879-01 SunOS 5.8: Solaris Product Registry patch SUNWwsr

Patches Applied To OS Services
Os Service Patch
-----
Solaris_8

Patches Applied To Clone Areas
Clone Area Patch
-----
Solaris_8/sun4u

```

Troubleshooting Diskless Client Problems

This section describes problems that are encountered when managing diskless clients and possible solutions.

Troubleshooting Diskless Client Installation Problems

The `smosservice add` command does not install any packages that are designated `ARCH=all` in the root (`/`) or `/usr` file systems. As a result, these packages are skipped. No warning or error messages are displayed. You must add these packages to the newly created Oracle Solaris OS service manually. This behavior has existed since the Solaris 2.1 OS. The behavior applies to both SPARC based and x86 based clients. Note that the list of missing packages varies, depending on which Oracle Solaris release you are running.

▼ How to Locate and Install Missing `ARCH=all` Packages

This procedure shows you how to locate and install missing `ARCH=all` packages after you have created the Oracle Solaris OS service on the server. Examples that are provided in this procedure apply to the Solaris 10 6/06 OS.

1 Locate all the packages with the `ARCH=all` parameter.

- a. Change directories to the **Product** directory of the media for the Oracle Solaris 10 image. For example:

```
% cd /net/server/export/Solaris/s10u2/combined.s10s_u2wos/latest/Solaris_10/Product
```

- b. List all the packages in the `pkginfo` file that have the `ARCH=all` parameter.

```
% grep -w ARCH=all */pkginfo
```

If an error message indicating the arguments list is too long is displayed, you can alternately run the following command to generate the list:

```
% find . -name pkginfo -exec grep -w ARCH=all {} /dev/null \;
```

Note that running this command takes longer to produce results.

The output is similar to the following:

```
./SUNWjdmk-base/pkginfo:ARCH=all
./SUNWjhdev/pkginfo:ARCH=all
./SUNWjhrt/pkginfo:ARCH=all
./SUNWjhdem/pkginfo:ARCH=all
./SUNWjhdoc/pkginfo:ARCH=all
./SUNWmlibk/pkginfo:ARCH=all
```

The information that is provided in this list enables you to determine which packages are installed in the `/usr` file system and which packages are installed in the root (`/`) file system.

c. Check the value of the SUNW_PKGTYPE parameter in the package list you generated.

Packages that belong in the /usr file system are designated as SUNW_PKGTYPE=usr in the pkginfo file. Packages that belong in the root (/) file system are designated as SUNW_PKGTYPE=root in the pkginfo file. In the preceding output, all the packages belong in the /usr file system.

2 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

3 Create the temporary installation administration files.

You must create a separate installation administration file for packages that are installed in the root (/) file system and a separate installation administration file for packages that are installed in the /usr file system.

- For ARCH=all packages that are installed in the /usr file system, create the following temporary installation administration file:

```
# cat >/tmp/admin_usr <<EOF
mail=
instance=unique
partial=nocheck
runlevel=nocheck
idepend=nocheck
rdepend=nocheck
space=nocheck
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=/usr_sparc.all
EOF
#
```

- For ARCH=all packages that are installed in the root (/) file system, if any exist, create the following temporary installation administration file:

```
# cat >/tmp/admin_root <<EOF
mail=
instance=unique
partial=nocheck
runlevel=nocheck
idepend=nocheck
rdepend=nocheck
space=nocheck
setuid=nocheck
conflict=nocheck
action=nocheck
EOF
#
```

4 Install the missing ARCH=all packages.

- a. If the current directory is not the Product directory of the media for the Oracle Solaris 10 image, change directories to that directory. For example:

```
# cd /net/server/export/Solaris/s10u2/combined.s10s_u2wos/latest/Solaris_10/Product
```

You can run the `pwd` command to determine the current directory.

- b. Install the missing ARCH=all packages in the `/usr` file system.

```
# pkgadd -R /export/Solaris_10 -a /tmp/admin_usr -d 'pwd' [
package-A package-B ...]
```

Multiple packages can be listed when you running the `pkgadd` command.

- c. Check that the ARCH=all packages were installed.

```
# pkginfo -R /export/Solaris_10 [package-A
package-B ...]
```

- d. Install the missing ARCH=all packages that in the root (`/`) file system.

Note that it is possible that none of these packages exist.

```
# pkgadd -R /export/root/clone/Solaris_10/sun4u -a /tmp/admin_root -d 'pwd' [
package-X package-Y ...]
```

- e. Check that the ARCH=all packages were installed.

```
# pkginfo -R /export/root/clone/Solaris_10/sun4u [
package-X package- ...]
```

- 5 After you have finished adding the missing ARCH=all packages, remove the temporary installation administration file.**

```
# rm /tmp/administration-file
```

Example 7-10 Locating and Installing Missing ARCH=all Packages

This example shows how to install the missing ARCH=all package, `SUNWjdmk-base`, in the `/usr` file system.

```
% uname -a
SunOS t1fac46 5.10 Generic_118833-02 sun4u sparc SUNW,UltraSPARC-IIi-cEngine

% cat /etc/release
      Oracle Solaris 10 8/11 s10x_u10wos_08 X86
  Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
        Assembled 28 February 2011

% cd /net/ventor/export/Solaris/s10u2/combined.s10s_u2wos/latest/Solaris_10/Product
% grep -w ARCH=all */pkginfo
Arguments too long

% find . -name pkginfo -exec grep -w ARCH=all {} /dev/null \;
./SUNWjdmk-base/pkginfo:ARCH=all
```

```

./SUNWjhdev/pkginfo:ARCH=all
./SUNWjhrt/pkginfo:ARCH=all
./SUNWjhdem/pkginfo:ARCH=all
./SUNWjhdoc/pkginfo:ARCH=all
./SUNWmllib/pkginfo:ARCH=all

% grep -w SUNW_PKGTYPE=usr ./SUNWjdmk-base/pkginfo ./SUNWjhdev/pkginfo ...
./SUNWjdmk-base/pkginfo:SUNW_PKGTYPE=usr
./SUNWjhdev/pkginfo:SUNW_PKGTYPE=usr
./SUNWjhrt/pkginfo:SUNW_PKGTYPE=usr
./SUNWjhdem/pkginfo:SUNW_PKGTYPE=usr
./SUNWjhdoc/pkginfo:SUNW_PKGTYPE=usr

% grep -w SUNW_PKGTYPE=root ./SUNWjdmk-base/pkginfo ./SUNWjhdev/pkginfo ...
% su
Password: xxxxxx
# cat >/tmp/admin_usr <<EOFmail=
instance=unique
partial=nocheck
runlevel=nocheck
idepend=nocheck k
rdepend=nocheck
space=nocheck
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=/usr_sparc.all
EOF

# pwd
/net/ventor/export/Solaris/s10u2/combined.s10s_u2wos/latest/Solaris_10/Product

# pkginfo -R /export/Solaris_10 SUNWjdmk-base
ERROR: information for "SUNWjdmk-base" was not found

# pkgadd -R /export/Solaris_10 -a /tmp/admin_usr -d 'pwd' SUNWjdmk-base

Processing package instance <SUNWjdmk-base> </net/ventor/export/Solaris/s10u2/combined.s10s_u2wos...

Java DMK 5.1 minimal subset(all) 5.1,REV=34.20060120
Copyright 2005 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Using </export/Solaris_10/usr_sparc.all>
## Processing package information.
## Processing system information.

Installing Java DMK 5.1 minimal subset as <SUNWjdmk-base>

## Installing part 1 of 1.
2438 blocks

Installation of <SUNWjdmk-base> was successful.

# pkginfo -R /export/Solaris_10 SUNWjdmk-base
application SUNWjdmk-base Java DMK 5.1 minimal subset

# rm /tmp/admin_usr

```

Troubleshooting General Diskless Client Problems

This section lists some common problems with diskless clients that you might encounter and possible solutions.

Problem: Diskless client reports Owner of the module `/usr/lib/security/pam_unix_session.so.1` is not root, when attempting to log in, the `/usr` file system is owned by nobody.

Solution: To correct the problem, follow this workaround:

1. Using a text editor, modify the diskless client's `server:/export/root/client/etc/default/nfs` file.
2. Change the `#NFMAPID_DOMAIN=domain` line to the following:
`NFMAPID_DOMAIN=the_same_value_as_in_server's_/var/run/nfs4_domain`
3. Ensure that the OS server and the diskless client have the same `nfsmapid` domain. To verify this information, check the `/var/run/nfs4_domain` file.



Caution – If the diskless client's `nfs4_domain` file contains a different value than the OS server's `/var/run/nfs4_domain` file, you will not be able to log in to the system after the diskless client boots.

4. Reboot the diskless client.

For more information, see [Chapter 3, “NFS Tunable Parameters,”](#) in *Oracle Solaris Tunable Parameters Reference Manual* and `nfsmapid(1M)`.

Problem: The OS server fails to do the following:

- Respond to client Reverse Address Resolution Protocol (RARP) requests
- Respond to client `bootparam` requests
- Mount a diskless client root (`/`) file system

Solution: The following solutions apply in a files environment.

- Verify that `files` is listed as the first source for `hosts`, `ethers`, and `bootparams` in the `/etc/nsswitch.conf` file on the OS server.
- Verify that the client's IP address appears in the `/etc/inet/hosts` file.

Note – If you are *not* running at least the Solaris 10 8/07 release, you must also verify that the client's IP address appears in the `/etc/inet/ipnodes` file.

In this Oracle Solaris release, the `/etc/inet/hosts` file is a single file that contains both IPv4 and IPv6 entries. You do not need to maintain IPv4 entries in two `hosts` files that always require synchronization. For backward compatibility, the `/etc/inet/ipnodes` file is replaced with a symbolic link with the same name to the `/etc/inet/hosts` file. For more information, see the [hosts\(4\)](#) man page.

- Verify that the client's Ethernet address appears in the `/etc/ethers` file.
- Verify that the `/etc/bootparams` file contains the following paths to the client's root (`/`) directory and swap areas.

```
client root=os-server:/export/root/
client swap=os-server:
/export/swap/client
```

The swap size varies depending on whether you specify the `-x swappsize` option when you add the diskless client. If you specify the `-x dump` option when you add the diskless client, the following line is present.

```
dump=os-server:/export/dump/client
dumpsize=512
```

The dump size varies depending on whether you specify the `-x dumpsize` option when you add the diskless client.

- Verify that the OS server's IP address appears in the `/export/root/client/etc/inet/hosts` file.

Problem: The OS server fails to do the following:

- Respond to client RARP requests
- Respond to client bootparam requests
- Mount a diskless client root (`/`) file system

Solution: The following solutions apply in a name service environment.

- Verify that both the OS server's and the client's Ethernet address and IP address are correctly mapped.
- Verify that the `/etc/bootparams` file contains the paths to the client's root (`/`) directory and swap areas.

```
client root=os-server:/export/
root/client swap=os-server:/export/
swap/client swappsize=24
```


The swap size varies depending on whether you specify the `-x swapsize` option when you add the diskless client. If you specify the `-x dump` option when you add the diskless client, the following line is present:

```
dump=os-server:/export/dump/
client dumpsize=24
```

The dump size varies depending on whether you specify the `-x dumpsize` option when you add the diskless client.

Problem: Diskless client panics

Solution: Verify the following:

- The OS server's Ethernet address is correctly mapped to its IP address. If you physically moved a system from one network to another, you might have forgotten to remap the system's new IP address.
- The client's host name, IP address, and Ethernet address do not exist in the database of another server *on the same subnet* that responds to the client's RARP, Trivial File Transfer Protocol (TFTP), or bootparam requests. Often, test systems are set up to install their OS from an install server. In these cases, the install server answers the client's RARP or bootparam request, returning an incorrect IP address. This incorrect address might result in the download of a boot program for the wrong architecture, or a failure to mount the client's root (/) file system.
- The diskless client's TFTP requests are not answered by an install server (or previous OS server) that transfers an incorrect boot program. If the boot program is of a different architecture, the client immediately panics. If the boot program loads from a non-OS server, the client might obtain its root partition from the non-OS server and its /usr partition from the OS server. In this situation, the client panics if the root and /usr partitions are of conflicting architectures or versions.
- If you are using both an install server and an OS server, verify that the following entry exists in the /etc/dfs/dfstab file:

```
share -F nfs -o -ro
/export/exec/Solaris_version- \
instruction-set.all/usr
```

where *version*= 8, 9, 10, and *instruction-set*=sparc or i386.

- Verify that the diskless client's root (/), /swap, and /dump (if specified) partitions have share entries:

```
share -F nfs -o rw=client,root=client
/export/root/client
share -F nfs -o rw=client,root=client /export/swap/
client
share -F nfs -o rw=client,root=client /export/dump/
client
```

- On the OS server, type the following command to check which files are shared:

% **share**

The OS server must share `/export/root/client` and `/export/swap/client-name` (defaults), or the root, `/swap`, and `/dump` partitions that you specified when you added the diskless client.

Verify that the following entries exist in the `/etc/dfs/dfsstab` file:

```
share -F nfs -o ro /export/exec/Solaris_version-
instruction-set.all/usr
share -F nfs -o rw=client,root=client /export/root/
client
share -F nfs -o rw=client,root=client /export/swap/
client
```

Problem: OS server is not responding to diskless client's RARP request

Solution: From the client's intended OS server, run the `snoop` command as superuser (`root`) by using the client's Ethernet address:

```
# snoop xx:xx:xx:xx:xx:xx
```

Problem: Boot program downloads but panics early in the process

Solution: Use the `snoop` command to verify that the intended OS server is answering the client's TFTP and NFS requests.

Problem: Diskless client hangs.

Solution: Restart the following daemons on the OS server:

```
# /usr/sbin/rpc.bootparamd
# /usr/sbin/in.rarpd -a
```

Problem: Incorrect server responds to diskless client's RARP request

Solution: Restart the following daemons on the OS server:

```
# /usr/sbin/rpc.bootparamd
# svcadm enable network/rarp
```

Introduction to Shutting Down and Booting a System

Oracle Solaris is designed to run continuously, so that electronic mail and network resources are available to users. This chapter provides guidelines for shutting down and booting a system.

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

- “What's New in Shutting Down and Booting a System” on page 171
- “Where to Find Shut Down and Boot Tasks” on page 176
- “Shut Down and Boot Terminology” on page 177
- “Guidelines for Shutting Down a System” on page 177
- “Guidelines for Booting a System” on page 178
- “When to Shut Down a System” on page 179
- “When to Boot a System” on page 180

For an overview of all of the boot features and methods that are available in the Oracle Solaris release, see [Chapter 9, “Shutting Down and Booting a System \(Overview\)”](#).

For instructions on booting a system, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\)”](#).

What's New in Shutting Down and Booting a System

This section describes new boot features in the Oracle Solaris release. For a complete listing of new features and a description of Oracle Solaris releases, see [Oracle Solaris 10 8/11 What's New](#). The following is a list of new features:

- “Support for Fast Reboot on the SPARC Platform” on page 172
- “Oracle Solaris Auto Registration Introduced” on page 172
- “Automatic Boot Archive Recovery” on page 173
- “SPARC Support for Install-Time Updates” on page 173
- “Two-Terabyte Disk Support for Installing and Booting Oracle Solaris 10” on page 174
- “Oracle Solaris ZFS Boot Support” on page 174

- “x86: findroot Command” on page 174
- “Support for Specifying Platform by Using bootadm Command” on page 175
- “Redesign of SPARC Bootstrap Process” on page 175
- “x86: Support for Using Power Button to Initiate System Shutdown” on page 176

Support for Fast Reboot on the SPARC Platform

The Fast Reboot feature of Oracle Solaris is now supported on the SPARC platform. The integration of Fast Reboot on the SPARC platform enables the `-f` option to be used with the `reboot` command to accelerate the boot process by skipping certain POST tests.

The Fast Reboot feature of Oracle Solaris is managed through SMF and implemented through a boot configuration service, `svc:/system/boot-config`. The `boot-config` service provides a means for setting or changing the default boot configuration parameters. When the `config/fastreboot_default` property is set to `true`, the system performs a fast reboot automatically, without the need to use the `reboot -f` command. By default, this property value is set to `false` on the SPARC platform.

Note – On the SPARC platform, the `boot-config` service also requires the `solaris.system.shutdown` authorization as the `action_authorization` and `value_authorization`.

To make Fast Reboot the default behavior on the SPARC platform, use the `svccfg` and `svcadm` commands.

For task-related information, see “[Accelerating the Reboot Process on the SPARC Platform \(Task Map\)](#)” on page 264.

Oracle Solaris Auto Registration Introduced

Oracle Solaris 10 9/10: Oracle Solaris Auto Registration feature is a mechanism whereby newly installed software products are automatically registered with My Oracle Support during the first system reboot after an installation or upgrade, and on subsequent system reboots, if any system configuration changes have occurred. Auto Registration leverages the existing service tag technology that enables products to be discovered on a network and then registered in a local registry.

The Auto Registration feature is managed by an SMF service. This service, which is enabled by default, runs once at boot time to check for newly installed products. If any new products are discovered, service tag information about these products is automatically transmitted to the Oracle Product Registration System by using an Hypertext Transfer Protocol Secure (HTTPS) connection.

The feature has a command-line interface (CLI), `usr/sbin/regadm`, that can be used by a privileged system administrator to manage the Auto Registration SMF service and to administer product registration, independent of the installation or upgrade process.

For more information, see [Chapter 17, “Working With the Oracle Solaris Auto Registration `regadm` Command \(Tasks\).”](#)

Automatic Boot Archive Recovery

Oracle Solaris 10 9/10: Starting with this release, boot archive recovery on the SPARC platform is automatic.

To support auto-recovery of the boot archives on the x86 platform, a new `auto-reboot-safe` property has been added to the boot configuration service, `svc:/system/boot-config:default`. By default, the property's value is set to `false` to ensure that the system does not automatically reboot to an unknown boot device. If the system is configured to automatically point to the BIOS boot device and GRUB menu entry that Oracle Solaris 10 is installed on, you can set the property's value to `true`. Setting the value to `true` enables an automatic reboot of the system for the purpose of recovering an out-of-date boot archive.

To set or change this property's value, use the `svccfg` and `svcadm` commands. See the [`svccfg\(1M\)`](#) and [`svcadm\(1M\)`](#) man pages.

For general information about this enhancement, see the [`boot\(1M\)`](#) man page.

For step-by-step instructions, see “[x86: How to Clear Automatic Boot Archive Update Failures by Using the `auto-reboot-safe` Property](#)” on page 270.

SPARC Support for Install-Time Updates

Oracle Solaris 10 9/10: Starting with this release, the `itu` utility has been modified to support booting a SPARC based system with Install-Time Updates (ITUs). Third-party vendors can now deliver driver updates on floppy disk, CD or DVD, and USB storage. In addition, new tools that enable you to modify the Oracle Solaris installation media with new packages and patches have been introduced. These tools can be used to deliver software updates for hardware platforms and to produce customized installation media. For task-related information, see “[SPARC: How to Boot a System With a Newly Created ITU](#)” on page 211.

See also the following man pages:

- [`itu\(1M\)`](#)
- [`mkbootmedia\(1M\)`](#)
- [`pkg2du\(1M\)`](#)

- [updatemedia\(1M\)](#)

Two-Terabyte Disk Support for Installing and Booting Oracle Solaris 10

Solaris 10 10/09: In previous releases, you could not install and boot the Solaris OS from a disk that was greater than 1 Tbyte in size. Starting with this release, you can install and boot the Oracle Solaris OS from a disk that is up to 2 Tbytes in size. In previous releases, you also had to use an EFI label for a disk that was larger than 1 Tbyte. In this release, you can use the VTOC label on any size disk. However, the addressable space by the VTOC label is limited to 2 Tbytes.

For more information, see [“What’s New in Disk Management?”](#) in *System Administration Guide: Devices and File Systems*.

Oracle Solaris ZFS Boot Support

Solaris 10 10/08: This release includes Oracle Solaris ZFS installation, as well as ZFS boot support. You can now install and boot from a ZFS root file system. This enhancement applies to both the SPARC and x86 based platforms. Booting, system operations, and installation procedures have been modified to support this change.

For more information, see [“Booting From an Oracle Solaris ZFS Root File System”](#) on page 190.

x86: findroot Command

All Oracle Solaris installation methods, including Solaris Live Upgrade, now use the `findroot` command for specifying which disk slice on an x86 based system to boot. This implementation supports booting systems with Oracle Solaris ZFS roots, as well as UFS roots. Previously, the `root` command, `root (hd0.0.a)`, was used to explicitly specify which disk slice to boot. This information is located in the `menu.lst` file that is used by GRUB.

The most common form of the GRUB `menu.lst` entry is now as follows:

```
findroot (rootfs0,0,a)
kernel$ /platform/i86pc/kernel/$ISADIR/unix -B $ZFS-BOOTFS
module$ /platform/i86pc/$ISADIR/boot_archive
```

In some Oracle Solaris 10 releases, the entry is as follows:

```
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive
```

For more information, see [“x86: Implementation of the findroot Command”](#) on page 223.

Support for Specifying Platform by Using `bootadm` Command

A new `-p` option has been added to the `bootadm` command.

This option enables you to specify the platform or machine hardware class of a client system in situations where the client platform differs from the server platform, for example when administering diskless clients.

Note – The `-p` option must be used with the `-R` option.

```
# bootadm -p platform -R [altroot]
```

The specified platform must be one of the following:

- `i86pc`
- `sun4u`
- `sun4v`

For more information, see the [bootadm\(1M\)](#) man page.

Redesign of SPARC Bootstrap Process

The Oracle Solaris SPARC bootstrap process has been redesigned to increase commonality with the x86 boot architecture.

Other enhancements include an improved boot architecture that supports booting a system from additional file system types, for example an Oracle Solaris ZFS file system, or a single miniroot, for installation, as well as booting from DVD, NFS, or HTTP. These enhancements increase flexibility and reduce maintenance requirements on SPARC based systems.

As part of this redesign, the boot archives and the `bootadm` command, previously only available on the x86 based platform, are now an integral part of the SPARC boot architecture.

The primary difference between the SPARC and x86 boot architectures is how the boot device and file are selected at boot time. The SPARC based platform continues to use the OpenBoot PROM (OBP) as the primary administrative interface, with boot options selected by using OBP commands. On x86 based systems, these options are selected through the BIOS and the GRand Unified Bootloader (GRUB) menu.

Note – Although the SPARC boot process has changed, no administrative procedures for booting a SPARC based system have been impacted. Boot tasks performed by the system administrator remain the same as they were prior to the boot architecture redesign.

For more information, see the [boot\(1M\)](#) and [bootadm\(1M\)](#) man pages.

For more information in this document, see “[Understanding the New SPARC Boot Architecture](#)” on page 185.

x86: Support for Using Power Button to Initiate System Shutdown

Pressing and releasing the power button on x86 based systems initiates a clean system shutdown and turns the system off. This functionality is equivalent to using the `init 5` command to shut down a system. On some x86 based systems, the BIOS configuration might prevent the power button from initiating shutdown. To enable use of the power button to perform a clean system shutdown, reconfigure the BIOS.

Note – On certain x86 based systems that were manufactured before 1999 and are running an older release, pressing the power button immediately turns off system power without safely shutting it down. This same behavior occurs when you press the power button on systems running with ACPI support that is disabled through the use of `acpi-user-options`.

For more information about `acpi-user-options`, see the [eeprom\(1M\)](#) man page.

Where to Find Shut Down and Boot Tasks

Use these references to find step-by-step instructions for shutting down and booting a system.

Shut Down and Boot Task	For More Information
Shut down a SPARC or an x86 based system	Chapter 10, “Shutting Down a System (Tasks)”
Modify boot behavior	Chapter 11, “Modifying Oracle Solaris Boot Behavior (Tasks)”
Boot a SPARC based system or an x86 based system	Chapter 12, “Booting an Oracle Solaris System (Tasks)”
Manage the Solaris boot archives	Chapter 13, “Managing the Oracle Solaris Boot Archives (Tasks)”

Shut Down and Boot Task	For More Information
Troubleshoot boot behavior on a SPARC or an x86 based system	“ Troubleshooting Booting on the SPARC Platform (Task Map) ” on page 277

Shut Down and Boot Terminology

The following terminology is used when shutting down and booting a system:

Run levels and init states A *run level* is a letter or digit that represents a system state in which a particular set of system services are available. The system is always running in one of a set of well-defined run levels. Run levels are also referred to as *init states* because the `init` process maintains the run level. System administrators use the `init` command or the `svcadm` command to initiate a run-level transition. This book refers to init states as run levels.

Boot options A *boot option* describes how a system is booted.

Different boot options include the following:

- **Interactive boot** – You are prompted to provide information about how the system is booted, such as the kernel and device path name.
- **Reconfiguration boot** – The system is reconfigured to support newly added hardware or new pseudo devices.
- **Recovery boot** – The system is hung or an invalid entry is prohibiting the system from booting successfully or from allowing users to log in.

For terminology that is specific to GRUB based booting, see “[x86: GRUB Terminology](#)” on page 290.

Guidelines for Shutting Down a System

Keep the following in mind when you shut down a system:

- Use the `init` and `shutdown` commands to shut down a system. Both commands perform a clean system shutdown, which means that all system processes and services are terminated normally.

x86 only – For x86 based systems that are running at least the Solaris 10 6/06 release, you can initiate a clean system shutdown by pressing and releasing the power button. Shutting down an x86 based system in this manner is equivalent to using the `init 5` command to shut down a system. On some x86 based systems, the BIOS configuration might prevent the power button from initiating a system shutdown. To use the power button, reconfigure the BIOS.

- Use the `shutdown` command to shut down a server. Logged-in users and systems that mount resources from the server are notified before the server is shut down. Additional notification of system shutdowns by electronic mail is also recommended so that users can prepare for system downtime.
- You need superuser privileges to use the `shutdown` or `init` command to shut down a system.
- Both `shutdown` and `init` commands take a run level as an argument.

The three most common run levels are as follows:

- **Run level 3** – All system resources are available and users can log in. By default, booting a system brings it to run level 3, which is used for normal day-to-day operations. This run level is also known as multiuser level with NFS resources shared.
- **Run level 6** – Stops the operating system and reboots to the state that is defined by the `initdefault` entry in the `/etc/inittab` file.
- **Run level 0** – The operating system is shut down, and it is safe to turn off power. You need to bring a system to run level 0 whenever you move a system, or add or remove hardware.

Run levels are fully described in [Chapter 18, “Managing Services \(Overview\)”](#).

Guidelines for Booting a System

Keep the following in mind when you boot a system:

- After a SPARC based system is shut down, it is booted by using the `boot` command at the PROM level.
- After an x86 based system is shut down, it is booted by selecting an OS instance in the GRUB menu.
- In the Solaris 9 release and some Oracle Solaris 10 releases, after an x86 based system is shut down, it is booted by using the `boot` command at the Primary Boot Subsystem menu.
- A system can be rebooted by turning the power off and then back on.



Caution – This method is not considered a clean shutdown, unless you have an x86 based system that is running a release that supports this shutdown method. See “[x86: Support for Using Power Button to Initiate System Shutdown](#)” on page 176. Use this shutdown method only as an alternative in emergency situations. Because system services and processes are terminated abruptly, file system damage is likely to occur. The work required to repair this type of damage could be substantial and might require the restoration of various user and system files from backup copies.

- SPARC and x86 based systems use different hardware components for booting. These differences are described in [Chapter 15, “x86: GRUB Based Booting \(Reference\).”](#)

When to Shut Down a System

The following table lists system administration tasks and the type of shutdown method that is required to initiate the task.

TABLE 8-1 Shutting Down a System

Reason for System Shutdown	Appropriate Run Level	For More Information
To turn off system power due to anticipated power outage.	Run level 0, where it is safe to turn off power	Chapter 10, “Shutting Down a System (Tasks)”
To change kernel parameters in the <code>/etc/system</code> file.	Run level 6 (reboot the system)	Chapter 10, “Shutting Down a System (Tasks)”
To perform file system maintenance, such as backing up or restoring system data.	Run level S (single-user level)	Chapter 10, “Shutting Down a System (Tasks)”
To repair a system configuration file such as <code>/etc/system</code> .	See “ When to Boot a System ” on page 180	N/A
To add or remove hardware from the system.	Reconfiguration boot (also to turn off power when adding or removing hardware)	“ Adding a Peripheral Device to a System ” in <i>System Administration Guide: Devices and File Systems</i>
To repair an important system file that is causing system boot failure.	See “ When to Boot a System ” on page 180	N/A
To boot the kernel debugger (<code>kmdb</code>) to track down a system problem.	Run level 0, if possible	Chapter 10, “Shutting Down a System (Tasks)”
To recover from a hung system and force a crash dump.	See “ When to Boot a System ” on page 180	N/A

TABLE 8-1 Shutting Down a System (Continued)

Reason for System Shutdown	Appropriate Run Level	For More Information
Reboot the system by using the kernel debugger (kldb), if the debugger can't be loaded at runtime.	Run level 6 (reboot the system)	<p>“SPARC: How to Boot the System With the Kernel Debugger (kldb)” on page 283</p> <p>“x86: How to Boot a System With the Kernel Debugger in the GRUB Boot Environment (kldb)” on page 286</p>

For examples of shutting down a server or a stand-alone system, see Chapter 10, “Shutting Down a System (Tasks).”

When to Boot a System

The following table lists system administration tasks and the corresponding boot option that is used to complete the task.

TABLE 8-2 Booting a System

Reason for System Reboot	Appropriate Boot Option	Information for SPARC Based Systems	Information for x86 Based Systems
Turn off system power due to anticipated power outage.	Turn system power back on.	Chapter 10, “Shutting Down a System (Tasks)”	Chapter 10, “Shutting Down a System (Tasks)”
Change kernel parameters in the <code>/etc/system</code> file.	Reboot the system to run level 3 (multiuser level with NFS resources shared).	“SPARC: How to Boot a System to Run Level 3 (Multiuser Level)” on page 229	“x86: How to Boot a System to Run Level 3 (Multiuser)” on page 246
Perform file system maintenance, such as backing up or restoring system data.	Press Control-D from run level S to bring the system back to run level 3.	“SPARC: How to Boot a System to Run Level S (Single-User Level)” on page 230	“x86: How to Boot a System to Run Level S (Single-User Level)” on page 247
Repair a system configuration file such as <code>/etc/system</code> .	Interactive boot.	“SPARC: How to Boot a System Interactively” on page 231	“x86: How to Boot a System Interactively” on page 249
Add or remove hardware from the system.	Reconfiguration boot (also to turn on system power after adding or removing hardware).	“Setting Up Disks for UFS File Systems (Task Map)” in <i>System Administration Guide: Devices and File Systems</i>	“Setting Up Disks for UFS File Systems (Task Map)” in <i>System Administration Guide: Devices and File Systems</i>
Boot the system by using the kernel debugger (kldb) to track down a system problem.	Booting with the kldb option.	“SPARC: How to Boot the System With the Kernel Debugger (kldb)” on page 283	“x86: How to Boot a System With the Kernel Debugger in the GRUB Boot Environment (kldb)” on page 286

TABLE 8-2 Booting a System

(Continued)

Reason for System Reboot	Appropriate Boot Option	Information for SPARC Based Systems	Information for x86 Based Systems
Boot the system in failsafe mode to repair an important system file that is causing system boot failure.	Booting the failsafe archive.	“How to Boot a SPARC Based System in Failsafe Mode” on page 239	“How to Boot an x86 Based System in Failsafe Mode” on page 256
To recover from a hung system and force a crash dump.	Performing a recovery boot	“SPARC: How to Force a Crash Dump and Reboot of the System” on page 279	“x86: How to Force a Crash Dump and Reboot of the System” on page 284

Shutting Down and Booting a System (Overview)

This chapter provides an overview of booting a system. The Oracle Solaris boot design, boot processes, and various methods of booting a system in the Oracle Solaris OS are described.

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

- “Fundamentals of the Oracle Solaris Boot Design” on page 184
- “Understanding the New SPARC Boot Architecture” on page 185
- “Implementation of the Boot Archives on SPARC” on page 187
- “x86: Administering the GRUB Bootloader” on page 188
- “Booting From an Oracle Solaris ZFS Root File System” on page 190

For instructions on booting an Oracle Solaris system, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\)”](#)

For instructions on booting a Solaris system that does *not* implement GRUB, see [Chapter 16, “x86: Booting a System That Does Not Implement GRUB \(Tasks\)”](#).

For what's new in shutting down and booting a system, see [“What's New in Shutting Down and Booting a System”](#) on page 171.

For overview information and instructions on administering boot loaders and modifying boot behavior, see [Chapter 11, “Modifying Oracle Solaris Boot Behavior \(Tasks\)”](#).

For information about managing boot services through the Service Management Facility (SMF), see [“SMF and Booting”](#) on page 355.

Fundamentals of the Oracle Solaris Boot Design

Note – The information in this section applies to both the SPARC and x86 platforms.

The fundamental Oracle Solaris boot design includes the following characteristics:

- **Use of a boot archive**

The boot archive is a ramdisk image that contains all of the files that are required for booting a system. When you install the Solaris OS, two boot archives are created, one primary archive and one failsafe archive. For more information, see [“Implementation of the Boot Archives on SPARC” on page 187](#).

The `bootadm` command has also been modified for use on the SPARC platform. This command functions the same way that it does on the x86 platform. The `bootadm` command handles the details of archive update and verification automatically. During an installation or system upgrade, the `bootadm` command creates the initial boot archive. During the process of a normal system shutdown, the shutdown process checks the boot archive contents against the root file system. If there are any inconsistencies, the system rebuilds the boot archive to ensure that on reboot, the boot archive and root (`/`) file system are synchronized. You can also use the `bootadm` command to manually update the boot archives. See [“Using the bootadm Command to Manage the Boot Archives” on page 272](#).

Note – Some options of the `bootadm` command cannot be used on SPARC based systems.

For more information, see the [`bootadm\(1M\)`](#) and [`boot\(1M\)`](#) man pages.

- **Use of a ramdisk image as the root file system during installation and failsafe operations**

This process is now the same on the SPARC and x86 platforms. The ramdisk image is derived from the boot archive and is then transferred to the system from the boot device.

Note – On the SPARC platform, the OpenBoot PROM continues to be used to access the boot device and to transfer the boot archive to the system's memory. Conversely, on the x86 platform, the system is initially controlled by the BIOS. The BIOS is used to initiate a transfer of the boot archive from a network device or to run a boot loader. In the Oracle Solaris OS, the x86 boot loader that is used to transfer the boot archive from disk is GRUB. See [“x86: Boot Processes” on page 289](#).

In the case of a software installation, the ramdisk image is the root file system that is used for the entire installation process. Using the ramdisk image for this purpose eliminates the need to boot the system from removable media. The ramdisk file system type can be either a High Sierra File System (HSFS) or UFS.

Understanding the New SPARC Boot Architecture

The boot processes on the SPARC platform have been redesigned and improved to increase commonality with the x86 boot experience. The new SPARC boot design enables the addition of new features, for example new file system types, without necessitating any changes to multiple portions of the boot chain. Changes also include the implementation of boot phase independence.

Highlights of these improvements include:

- Commonality in boot processes on the SPARC and x86 platforms
- Commonality in the network boot experience
- Boot architecture flexibility that enables booting a system from different file system types more easily

The following four boot phases are now independent of each other:

1. Open Boot PROM (OBP) phase

The OBP phase of the boot process on the SPARC platform is unchanged.

For disk devices, the firmware driver usually uses the OBP label package's *load* method, which parses the VTOC label at the beginning of the disk to locate the specified partition. Sectors 1-15 of the partition are then read into the system's memory. This area is commonly called the boot block and usually contains a file system reader.

2. Booter phase

During this phase the boot archive is read and executed. Note that this is the only phase of the boot process that requires knowledge of the boot file system format. In some instances, the boot archive might also be the installation miniroot. Protocols that are used for the transfer of the boot loader and the boot archive include local disk access, NFS, and HTTP.

3. Ramdisk phase

The ramdisk is a boot archive that is comprised of kernel modules any other components that are required to boot an instance of the Oracle Solaris OS. or an installation miniroot.

The SPARC boot archive is identical to an x86 boot archive. The boot archive file system format is private. Therefore, knowledge of the file system type that is used during a system boot, for example an HSFS or a UFS file system, is not required by the booter or the kernel. The ramdisk extracts the kernel image from the boot archive and then executes it. To minimize the size of the ramdisk, in particular, the installation miniroot that resides in the system's memory, the contents of the miniroot are compressed. This compression is performed on a per-file level and is implemented within the individual file system. The `/usr/sbin/fiocompress` utility is then used to compress the file and mark the file as compressed.

Note – This utility has a private interface to the file compression file system, `dcfs`.

4. Kernel phase

The kernel phase is the final stage of the boot process. During this phase, the Oracle Solaris OS is initialized and a minimal root file system is mounted on the ramdisk that was constructed from the boot archive. In some environments, such as an installation, the ramdisk is used as the root (`/`) file system and remains mounted. If the boot archive is the installation miniroot, the OS continues executing the installation process. Otherwise, the ramdisk contains a set of kernel files and drivers that is sufficient to mount the root file system on the specified root device.

The kernel then extracts the remainder of the primary modules from the boot archive, initializes itself, mounts the real root file system, then discards the boot archive.

Packing and Unpacking the Miniroot

The ramdisk-based miniroot is packed and unpacked by the `root_archive` command. Note that only SPARC based systems that support the new boot architecture have the ability to pack and unpack a compressed version of the miniroot.



Caution – The Oracle Solaris 10 version of the `root_archive` tool is not compatible with versions of the tool that are included in other Oracle Solaris releases. Therefore, ramdisk manipulation should only be performed on a system that is running the same release as the archives.

For more information about packing and unpacking the miniroot, see the [root_archive\(1M\)](#) man page.

Software Installation and Upgrades

To install or upgrade the Oracle Solaris OS, you need to boot from either CD/DVD, or from the network. In both instances, the miniroot's root file system is the ramdisk. This process enables you to eject the Solaris boot CD or DVD without having to reboot the system. Note that the boot archive contains the entire miniroot. The construction of the installation DVD has been modified to use an HSFS boot block. The miniroot is then packed into a single UFS file that is loaded as the ramdisk. Note that the miniroot is used for all OS installation types.

Installation Memory Requirements

For Oracle Solaris 10 9/10, the minimum memory requirement to install a SPARC based system is 384 Mbytes of memory. This amount of memory enables a text-based installation *only*. For x86 based systems, the minimum memory requirement is 768 Mbytes of memory. Also, to run the installation GUI program a minimum of 768 Mbytes of memory is required.

Changes to the Network Boot Server Setup Process

The network boot server setup process has been modified. The boot server now serves a bootstrap program, as well as the ramdisk, which is downloaded and booted as a single miniroot for all installations, whether booting from CD/DVD, or performing a network installation by using NFS or HTTP. The administration of a network boot server for a network boot over both NFS or the wanboot program (HTTP) remains the same. However, the internal implementation of the network boot process has been modified as follows:

1. The boot server transfers a bootstrap in the form of a boot archive to the target system.
2. The target system unpacks the boot archive in a ramdisk.
3. The boot archive is then mounted as the initial read-only root device.

For more information about booting a SPARC based system, see [“Booting a SPARC Based System \(Task Map\)” on page 228](#).

Support for Booting Multiple Kernels

On SPARC based systems, when you boot the system from the ok prompt, the default boot device is automatically selected. An alternate boot device can be specified by changing the NVRAM variable for the boot - device. You can also specify an alternate boot device or an alternate kernel (boot file) from the command line at boot time. See [“SPARC: How to Boot a Kernel Other Than the Default Kernel” on page 232](#).

Implementation of the Boot Archives on SPARC

The boot archives, previously only available on the x86 platform, are now an integral part of the SPARC boot architecture.

The bootadm command has been modified for use on the SPARC platform. This command functions the same as it does on the x86 platform. The bootadm command handles the details of archive update and verification. On the x86 platform the bootadm command updates the GRUB menu during an installation or system upgrade. You can also use the bootadm command to manually manage the boot archives.

The boot archive service is managed by the Service Management Facility (SMF). The service instance for the boot archive is `svc:/system/boot-archive:default`. To enable, disable, or refresh this service use the `svcadm` command. For information about managing services by using SME, see [Chapter 18, “Managing Services \(Overview\)”](#).

On supported Solaris releases, for both SPARC and x86 based systems, there are two kinds of boot archives:

- Primary boot archive
- Failsafe boot archive

The files that are included in the SPARC boot archives are located in the `/platform` directory.

The contents of the `/platform` directory is divided into two groups of files:

- Files that are required for a sun4u boot archive
- Files that are required for sun4v boot archive

For information about managing the boot archives, see [“Managing the Oracle Solaris Boot Archives \(Task Map\)” on page 267](#).

x86: Administering the GRUB Bootloader

The open source GRand Unified Bootloader (GRUB) is the default boot loader on x86 based systems. GRUB is responsible for loading a boot archive into the system’s memory. A boot archive is a collection of critical files that is needed during system startup before the root file system is mounted. The boot archive is the interface that is used to boot the Oracle Solaris OS. You can find more information about GRUB at <http://www.gnu.org/software/grub/grub.html>. See also the `grub(5)` man page.

How GRUB Based Booting Works

After an x86 based system is powered on, the Basic Input/Output System (BIOS) initializes the CPU, the memory, and the platform hardware. When the initialization phase has completed, the BIOS loads the boot loader from the configured boot device and then transfers control of the system to the boot loader. The *boot loader* is the first software program that runs after you turn on a system. This program starts the boot process.

GRUB implements a menu interface that includes boot options that are predefined in a configuration file called the `menu.lst` file. GRUB also has a command-line interface that is accessible from the GUI menu interface that can be used to perform various boot functions, including modifying default boot behavior. In the Solaris OS, the GRUB implementation is compliant with the Multiboot Specification, which is described in detail at <http://www.gnu.org/software/grub/grub.html>.

Because the Oracle Solaris kernel is fully compliant with the Multiboot Specification, you can boot x86 based systems by using GRUB. With GRUB, you can boot various operating systems that are installed on a single x86 based system. For example, you can individually boot Oracle Solaris, Linux, or Windows by selecting the boot entry in the GRUB menu at boot time, or by configuring the menu .lst file to boot a specific OS by default.

Because GRUB is intuitive about file systems and kernel executable formats, you can load an operating system without recording the physical position of the kernel on the disk. With GRUB-based booting, the kernel is loaded by specifying its file name, and the drive and the partition where the kernel resides. For more information see [“Naming Conventions That Are Used for Configuring GRUB” on page 292.](#)

For step-by-step instructions on booting a system with GRUB, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245.](#)

See also the following man pages:

- [boot\(1M\)](#)
- [bootadm\(1M\)](#)
- [grub\(5\)](#)
- [installgrub\(1M\)](#)

GRUB Support for findroot Command

The `findroot` command, which functions similarly to the `root` command previously used by GRUB, has enhanced capabilities for discovering a targeted disk, regardless of the boot device. The `findroot` command also supports booting from an Oracle Solaris ZFS root file system.

The most common format for the menu .lst entry for this command is as follows:

```
findroot (rootfs0,0,a)
kernel$ /platform/i86pc/kernel/$ISADIR/unix
module$ /platform/i86pc/$ISADIR/boot_archive
```

In some Oracle Solaris release, the entry is as follows:

```
title Solaris 10 10/08 s10x_u6wos_03 X86
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title Solaris failsafe
findroot (pool_rpool,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```

For more information, see [“x86: Implementation of the findroot Command” on page 223.](#)

For GRUB reference information, see [Chapter 15, “x86: GRUB Based Booting \(Reference\)”](#)

Booting From an Oracle Solaris ZFS Root File System

Support for booting from an Oracle Solaris ZFS root file system has been added to Oracle Solaris. The installation software also includes support for system upgrades and patching of systems with ZFS roots. Booting, system operations, and installation procedures have been modified to support this change. Changes to booting include the implementation of a new boot architecture on the SPARC platform. The new SPARC boot design includes feature enhancements that increase commonality with the Solaris x86 boot architecture.

Before using this feature, check the [Oracle Solaris 10 8/11 Release Notes](#) to find out about any known issues.

For more information about Oracle Solaris ZFS, including a complete list of terms, see “ZFS Terminology” in [Oracle Solaris ZFS Administration Guide](#).

Installation Requirements for Oracle Solaris ZFS

Before performing a new installation of Oracle Solaris or using Oracle Solaris Live Upgrade to migrate a UFS root file system to an Oracle Solaris ZFS root file system, make sure the following requirements are met:

- **Solaris release information:**

The ability to install and boot from an Oracle Solaris ZFS root file system is available, starting with the Solaris 10 10/09 release. To perform an Oracle Solaris Live Upgrade operation to migrate to a ZFS root file system, you must have installed or upgraded to at least the Solaris 10 10/09 release.

- **Oracle Solaris ZFS storage pool space requirements:**

Because swap and dump devices are not shared in a ZFS root environment, the minimum amount of available pool space that is required for a bootable ZFS root file system is larger than for a bootable UFS root file system.

Swap volume size is calculated at half the size of physical memory, but no more than 2 Gbytes, and no less than 512 Mbytes. Dump volume size is calculated by the kernel, based on `dumppadm` information and the size of physical memory. You can adjust the size of your swap and dump volumes to sizes of your choosing either in an Oracle Solaris JumpStart profile or during an initial installation, as long as the new sizes support system operation. For more information, see “ZFS Support for Swap and Dump Devices” in [Oracle Solaris ZFS Administration Guide](#).

How Booting From an Oracle Solaris ZFS Root File System Works

Booting from an Oracle Solaris ZFS root file system works differently than booting from a UFS file system. Because ZFS applies several new concepts for installation and booting, some basic administrative practices for booting a system have changed. The most significant difference between booting from a ZFS root file system and booting from a UFS root file system is that with ZFS a device identifier does *not* uniquely identify a root file system, and thus a BE. With ZFS, a device identifier uniquely identifies a *storage pool*. A storage pool can contain multiple bootable datasets (root file systems). Therefore, in addition to specifying a boot device, a root file system within the pool that is identified by the boot device must also be specified.

On an x86 based system, if the boot device identified by GRUB contains a ZFS storage pool, the menu .lst file that is used to create the GRUB menu is located in the dataset at the root of that pool's dataset hierarchy. This dataset has the same name as the pool. There is one such dataset in each pool.

A *default bootable dataset* is the bootable dataset for the pool that is mounted at boot time and is defined by the root pool's `boot fs` property. When a device in a root pool is booted, the dataset that is specified by this property is then mounted as the root file system.

The new `boot fs pool` property is a mechanism that is used by the system to specify the default bootable dataset for a given pool. When a device in a root pool is booted, the dataset that is mounted by default as the root file system is the one that is identified by the `boot fs pool` property.

On a SPARC based system, the default `boot fs pool` property is overridden by using the new `-Z dataset` option of the boot command.

On an x86 based system, the default `boot fs pool` property is overridden by selecting an alternate boot environment in the GRUB menu at boot time.

SPARC: Options That Support Booting From an Oracle Solaris ZFS Root File System

On the SPARC platform, the following two boot options are new:

- The `-L` option, which is used to print a list of all the available BEs on a system.

```
ok boot -L
```

Note – The `-L` option is run from the `ok` prompt. This option only presents the list of available BEs on the system. To boot the system, use the `-Z` boot option.

- The `-Z` option of the `boot` command enables you to specify a bootable dataset other than the default dataset that is specified by the `boot fs` pool property.

`ok boot -Z dataset`

The list of BEs that are displayed when you use the `-L` option on a device that has a ZFS boot loader reflect the `menu.lst` entries that are available on that particular system. Along with the list of available BEs, instructions for selecting a BE and using the `-Z` option to boot the system are also provided. The dataset specified by the `boot fs` value for the menu item is used for all subsequent files that are read by the booter, for example, the boot archive and various configuration files that are located in the `/etc` directory. This dataset is then mounted as the root file system.

For step-by-step instructions, see [“Booting From a Specified ZFS Root File System on a SPARC Based System”](#) on page 234.

x86: Options That Support Booting From a ZFS Root File System

On the x86 platform, a new GRUB keyword, `$ZFS-BOOTFS` has been introduced. When booting an x86 based system, if the root file system that corresponds with the GRUB menu entry is a ZFS dataset, the GRUB menu entry contains the `-B` option with the `$ZFS-BOOTFS` token, by default. If you install a release that supports a ZFS boot loader, the GRUB `menu.lst` file is updated with this information automatically. The default bootable dataset is identified by the `boot fs` property.

On x86 based systems that are running a release that supports a ZFS boot loader, this information is included in the GRUB `menu.lst` file.

For step-by-step instructions on booting a system from ZFS, see [“x86: Booting From a Specified ZFS Root File System on an x86 Based System”](#) on page 252.

Shutting Down a System (Tasks)

This chapter describes the procedures for shutting down systems.

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

- “Shutting Down the System (Task Map)” on page 193
- “Shutting Down the System” on page 194
- “Turning Off Power to All Devices” on page 201

For overview information about system run levels, see Chapter 18, “Managing Services (Overview).”

For information on the procedures associated with run levels and boot files, see “Shutting Down the System (Task Map)” on page 193.

Shutting Down the System (Task Map)

Task	Description	For Instructions
Determine who is logged in to a system.	Use the <code>who</code> command to determine who is logged in to a system.	“How to Determine Who Is Logged in to a System” on page 196
Shut down a server.	Use the <code>shutdown</code> command with the appropriate options to shut down a server.	“How to Shut Down a Server” on page 196
Shut down a stand-alone system.	Use the <code>init</code> command and indicate the appropriate run-level to shut down a stand-alone system.	“How to Shut Down a Stand-Alone System” on page 199

Task	Description	For Instructions
Turn off power to all devices.	Powering down a system includes the following devices: <ul style="list-style-type: none"> ■ CPU ■ Monitor ■ External devices, such as disks, tapes, and printers 	“How to Turn Off Power to All Devices” on page 201

Shutting Down the System

Oracle Solaris is designed to run continuously, so that the electronic mail and network software can work correctly. However, some system administration tasks and emergency situations require that the system be shut down to a level where it is safe to remove power. In some cases, the system needs to be brought to an intermediate level, where not all system services are available.

Such cases include the following:

- Adding or removing hardware
- Preparing for an expected power outage
- Performing file system maintenance, such as a backup

For a complete list of system administration tasks that require a system shutdown, see [Chapter 9, “Shutting Down and Booting a System \(Overview\)”](#).

For information on using your system's power management features, see the [pmconfig\(1M\)](#) man page.

System Shutdown Commands

The use of the `init` and `shutdown` commands are the primary ways to shut down a system. Both commands perform a *clean shutdown* of the system. As such, all file system changes are written to the disk, and all system services, processes, and the operating system are terminated normally.

The use of a system's Stop key sequence or turning a system off and then on are not clean shutdowns because system services are terminated abruptly. However, sometimes these actions are needed in emergency situations. For instructions on system recovery techniques, see [Chapter 12, “Bootting an Oracle Solaris System \(Tasks\)”](#), and [Chapter 13, “Managing the Oracle Solaris Boot Archives \(Tasks\)”](#).

Note – On x86 systems that are running at least the Solaris 10 6/06 release, pressing and releasing the power button initiates a clean system shutdown. This method is equivalent to using the `init 5` command.

The following table describes the various shutdown commands and provides recommendations for using them.

TABLE 10-1 Shutdown Commands

Command	Description	When To Use
<code>shutdown</code>	An executable shell script that calls the <code>init</code> program to shut down the system. The system is brought to run level <code>S</code> by default.	Recommended for servers operating at run level <code>3</code> because users are notified of the impending shutdown. Also notified are the systems that are mounting resources from the server that is being shut down.
<code>init</code>	An executable that kills all active processes and synchronizes the disks before changing run levels.	Recommended for stand-alone systems when other users will not be affected. Provides a faster system shutdown because users are not notified of the impending shutdown.
<code>reboot</code>	An executable that synchronizes the disks and passes boot instructions to the <code>uadmin</code> system call. In turn, this system call stops the processor.	The <code>init</code> command is the preferred method.
<code>halt</code> , <code>poweroff</code>	An executable that synchronizes the disks and stops the processor.	Not recommended because it doesn't shutdown all processes, and unmount any remaining file systems. Stopping the services, without doing a clean shutdown, should only be done in an emergency or if most of the services are already stopped.

User Notification of System Down Time

When the `shutdown` command is initiated, a warning followed by a final shutdown message is broadcast to all users who are currently logged in to the system and all systems that are mounting resources from the affected system.

For this reason, the `shutdown` command is preferred instead of the `init` command when you need to shut down a server. When you use either command, you might want to give users more notice by sending them a mail message about any scheduled system shutdown.

Use the `who` command to determine which users on the system need to be notified. This command is also useful for determining a system's current run level. For more information, see [“Determining a System's Run Level” on page 357](#) and the `who(1)` man page.

▼ How to Determine Who Is Logged in to a System

- 1 Log in to the system to be shut down.
- 2 Display all of the users who are logged in to the system.

```
$ who
```

Example 10-1 Determining Who Is Logged in to a System

The following example shows how to display who is logged in to the system.

```
$ who
holly      console    May  7 07:30
kryten     pts/0      May  7 07:35 (starlite)
lister     pts/1      May  7 07:40 (bluemidget)
```

- Data in the first column identifies the user name of the logged-in user.
- Data in the second column identifies the terminal line of the logged-in user.
- Data in the third column identifies the date and time that the user logged in.
- Data in the fourth column, if present, identifies the host name if a user is logged in from a remote system.

▼ How to Shut Down a Server

- 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

- 2 Find out if users are logged in to the system.

```
# who
```

A list of all logged-in users is displayed. You might want to send mail or broadcast a message to let users know that the system is being shut down.

- 3 Shut down the system.

```
# shutdown -iinit-level -ggrace-period -y
```

`-iinit-level` Brings the system to an init level that is different from the default of S. The choices are 0, 1, 2, 5, and 6.

Run levels 0 and 5 are reserved states for shutting the system down. Run level 6 reboots the system. Run level 2 is available as a multi-user operating state.

- ggrace-period* Indicates a time (in seconds) before the system is shut down. The default is 60 seconds.
- y Continues to shut down the system without intervention. Otherwise, you are prompted to continue the shutdown process after 60 seconds.

For more information, see the [shutdown\(1M\)](#) man page.

4 If you are asked for confirmation, type y.

Do you want to continue? (y or n): **y**

If you used the shutdown -y command, you will not be prompted to continue.

5 Type the superuser password, if prompted.

Type Ctrl-d to proceed with normal startup,
(or give root password for system maintenance): **xxxxxx**

6 After you have finished the system administration tasks, press Control-D to return to the default system run level.

7 Use the following table to verify that the system is at the run level that you specified in the shutdown command.

Specified Run Level	SPARC Based System Prompt	x86 Based System Prompt
S (single-user level)	#	#
0 (power-down level)	ok or >	Press any key to reboot
Run level 3 (multiuser level with remote resources shared)	<i>hostname</i> console login:	<i>hostname</i> console login:

Example 10–2 SPARC: Bringing a Server to Run Level S

In the following example, the shutdown command is used to bring a SPARC based system to run level S (single-user level) in three minutes.

```
# who
root  console      Jun 14 15:49    (:0)

# shutdown -g180 -y

Shutdown started.      Mon Jun 14 15:46:16 MDT 2004

Broadcast Message from root (pts/4) on venus Mon Jun 14 15:46:16...
The system venus will be shut down in 3 minutes .
.
.
Broadcast Message from root (pts/4) on venus Mon Jun 14 15:46:16...
The system venus will be shut down in 30 seconds .
```

```
.
.
INIT: New run level: S
The system is coming down for administration. Please wait.
Unmounting remote filesystems: /vol nfs done.
Shutting down Solaris Management Console server on port 898.
Print services stopped.
Jun 14 15:49:00 venus syslogd: going down on signal 15
Killing user processes: done.

Requesting System Maintenance Mode
SINGLE USER MODE

Root password for system maintenance (control-d to bypass): xxxxxx
single-user privilege assigned to /dev/console.
Entering System Maintenance Mode
#
```

Example 10-3 SPARC: Bringing a Server to Run Level 0

In the following example, the shutdown command is used to bring a SPARC based system to run level 0 in 5 minutes without requiring additional confirmation.

```
# who
root      console      Jun 17 12:39
userabc   pts/4              Jun 17 12:39  (:0.0)
# shutdown -i0 -g300 -y
Shutdown started.   Thu Jun 17 12:40:25 MST 2004

Broadcast Message from root (console) on pretend Thu Jun 17 12:40:25...
The system pretend will be shut down in 5 minutes
.
.
.
Changing to init state 0 - please wait
#
INIT: New run level: 0
The system is coming down. Please wait.
System services are now being stopped.
.
.
.
The system is down.
syncing file systems... done
Program terminated
Type help for more information
ok
```

If you are bringing the system to run level 0 to turn off power to all devices, see [“How to Turn Off Power to All Devices” on page 201](#).

Example 10-4 SPARC: Rebooting a Server to Run Level 3

In the following example, the shutdown command is used to reboot a SPARC based system to run level 3 in two minutes. No additional confirmation is required.

```

# who
root          console      Jun 14 15:49    (:0)
userabc      pts/4          Jun 14 15:46    (:0.0)
# shutdown -i6 -g120 -y
Shutdown started.   Mon Jun 14 15:46:16 MDT 2004

Broadcast Message from root (pts/4) on venus Mon Jun 14 15:46:16...
The system venus will be shut down in 2 minutes

Changing to init state 6 - please wait
#
INIT: New run level: 6
The system is coming down. Please wait.
.
.
.
The system is down.
syncing file systems... done
rebooting...
.
.
.
venus console login:

```

See Also Regardless of why you shut down a system, you will probably want to return to run level 3, where all file resources are available, and users can log in. For instructions on bringing a system back to a multiuser level, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\)”](#).

▼ How to Shut Down a Stand-Alone System

Use this procedure when you need to shut down a stand-alone system.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

2 Shut down the system.

```
# init 5
```

For more information, see the [init\(1M\)](#) man page.

- **Alternately, you can use the `uadmin` command to shut down the system.**

```
# uadmin 2 0
```

- **If you have an x86 based system that is running at least the Solaris 10 6/06 release, you can press and release the power button to initiate a clean system shutdown and turn off the system.**

This functionality is equivalent to using the `init 5` command to shut down a system. For more information, see “[What's New in Shutting Down and Booting a System](#)” on page 171.

- 3 Use the following table to verify that the system is at the run level that you specified in the `init` command.

Specified Run Level	SPARC Based System Prompt	x86 Based System Prompt
S (single-user level)	#	#
2 (multiuser level)	#	#
0 (power-down level)	ok or >	Press any key to reboot
3 (multiuser level with NFS resources shared)	<i>hostname</i> console login:	<i>hostname</i> console login:

Example 10-5 Using the `uadmin` command to Shut Down a System

```
# uadmin 2 0
syncing file systems... done
Program terminated
```

Example 10-6 Bringing a Stand-Alone System to Run Level 0

In this example, the `init` command is used to bring an x86 based stand-alone system to the level where it is safe to turn off power.

```
# init 0
#
INIT: New run level: 0
The system is coming down. Please wait.
.
.
.
The system is down.
syncing file systems... [11] [10] [3] done
Press any key to reboot
```

If you are bringing the system to run level 0 to turn off power to all devices, see [“How to Turn Off Power to All Devices”](#) on page 201.

Example 10-7 SPARC: Bringing a Stand-Alone System to Run Level S

In this example, the `init` command is used to bring a SPARC based stand-alone system to run level S (single-user level).

```
# init s
#
INIT: New run level: S
The system is coming down for administration. Please wait.
```



```

Unmounting remote filesystems: /vol nfs done.
Print services stopped.
syslogd: going down on signal 15
Killing user processes: done.

SINGLE USER MODE

Root password for system maintenance (control-d to bypass): xxxxxx
single-user privilege assigned to /dev/console.
Entering System Maintenance Mode
#

```

See Also Regardless of why you shut down the system, you will probably want to return to run level 3, where all file resources are available, and users can log in. For instructions on bringing a system back to a multiuser level, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\)”](#).

Turning Off Power to All Devices

You need to turn off power to all system devices when you do the following:

- Replace or add hardware.
- Move the system from one location to another.
- Prepare for an expected power outage or natural disaster such as an approaching electrical storm.

Turn off the power to system devices, including the CPU, the monitor, and external devices such as disks, tapes, and printers.

Before you turn off the power to all system devices, shut down the system cleanly, as described in the preceding sections.

▼ How to Turn Off Power to All Devices

- 1 Select one of the following methods to shut down the system:
 - If you are shutting down a server, see [“How to Shut Down a Server” on page 196](#).
 - If you are shutting down a stand-alone system, see [“How to Shut Down a Stand-Alone System” on page 199](#).
- 2 Turn off the power to all devices after the system is shutdown. If necessary, also unplug the power cables.

3 After power can be restored, use the following steps to turn on the system and devices.

- a. Plug in the power cables.
- b. Turn on the monitor.
- c. Turn on disk drives, tape drives, and printers.
- d. Turn on the CPU.

The system is brought to run level 3.

Modifying Oracle Solaris Boot Behavior (Tasks)

This chapter provides information about modifying boot behavior on Solaris systems.

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

- [“Modifying Boot Behavior on SPARC Based Systems \(Task Map\)”](#) on page 203
- [“Modifying Boot Behavior on x86 Based Systems \(Task Map\)”](#) on page 212

For what’s new in booting and general overview information about the boot process, see [Chapter 8, “Introduction to Shutting Down and Booting a System.”](#)

For step-by-step instructions on booting an Oracle Solaris system, see [Chapter 12, “Bootting an Oracle Solaris System \(Tasks\).”](#)

Modifying Boot Behavior on SPARC Based Systems (Task Map)

Task	Description	For Instructions
Identify the PROM revision number.	Use the banner command at the ok prompt to display the PROM revision number for a system.	“SPARC: How to Find the PROM Revision Number for a System” on page 205
Identify devices on the system that can be booted.	Before modifying boot behavior by using the boot PROM, identify the devices on the system.	“SPARC: How to Identify Devices on a System” on page 205
Display the current boot device.	Use this procedure to determine the current default boot device from which the system will boot.	“SPARC: How to Determine the Default Boot Device” on page 207

Task	Description	For Instructions
Change the default boot device.	To change the default boot device, use one of the following methods: <ul style="list-style-type: none"> ■ Change the boot - device parameter at the boot PROM. ■ Change the boot - device parameter by using the eeprom command. 	“SPARC: How to Change the Default Boot Device by Using the Boot PROM” on page 207 “SPARC: How to Change the Default Boot Device by Using the eeprom Command” on page 209
Reset the system.	When you reset the system, the system runs diagnostic tests on the hardware, then reboots.	“SPARC: Resetting the System” on page 209
Change the default boot file.	To change the default kernel that the system boots, use one of the following methods: <ul style="list-style-type: none"> ■ Change the boot - file parameter by using the boot PROM. ■ Change the boot - file parameter by using the eeprom command. 	“SPARC: How to Change the Default Kernel by Using the Boot PROM” on page 209 “SPARC: How to Change the Default Kernel by Using the eeprom Command” on page 210
Boot a system with install-time Updates (ITUs).	The new ITU utility enables you to boot a SPARC system with install-time updates that you create.	“SPARC: How to Boot a System With a Newly Created ITU” on page 211

SPARC: Using the Boot PROM

The boot PROM is used to boot a system. You might need to change the way the system boots. For example, you might want to reset the device to boot from or run hardware diagnostics before you bring the system to a multiuser level.

System administrators typically use the PROM level to boot a system. You can also change the default boot file and boot device at the PROM level.

If you need to perform any of the following tasks, you need to change the default boot device:

- Add a new drive to the system either permanently or temporarily
- Change the network boot strategy
- Temporarily boot a stand-alone system from the network

For a complete list of PROM commands, see the [monitor\(1M\)](#) and [eeprom\(1M\)](#) man pages.

▼ SPARC: How to Find the PROM Revision Number for a System

- Display a system's PROM revision number by using the **banner** command.

```
ok banner
Sun Ultra 5/10 UPA/PCI (UltraSPARC-IIi 333MHz), No Keyboard
OpenBoot 3.15, 128 MB memory installed, Serial #number.
Ethernet address number, Host ID: number.
```

Hardware configuration information, including the revision number of the PROM, is displayed. In the preceding example, the PROM revision number is 3.15.

▼ SPARC: How to Identify Devices on a System

You might need to identify the devices on the system to determine what are the appropriate devices to boot from.

Before You Begin Before you can safely use the probe commands to determine what devices are attached to the system, you need to do the following:

- Change the PROM auto-boot? parameter to false.
- Issue the reset-all command to clear system registers.

```
ok setenv auto-boot? false
```

```
ok reset-all
```

You can view the probe commands that are available on your system by using the **sifting** probe command:

```
ok sifting probe
```

If you run the probe commands without clearing the system registers, the following message is displayed:

```
ok probe-scsi
This command may hang the system if a Stop-A or halt command
has been executed. Please type reset-all to reset the system
before executing this command.
Do you wish to continue? (y/n) n
```

- 1 Identify the devices on the system.

```
ok probe-device
```

- 2 (Optional) If you want the system to reboot after a power failure or after using the reset command, then reset the `auto-boot?` parameter to true.

```
ok setenv auto-boot? true
auto-boot? = true
```

- 3 Boot the system to multiuser mode.

```
ok reset-all
```

Example 11-1 SPARC: Identifying the Devices on a System

The following example shows how to identify the devices connected to an Ultra 10 system.

```
ok setenv auto-boot? false
auto-boot? = false
ok reset-all
Resetting ...
```

```
Sun Ultra 5/10 UPA/PCI (UltraSPARC-IIi 333MHz), No Keyboard
OpenBoot 3.15, 128 MB memory installed, Serial #10933339.
Ethernet address 8:0:20:a6:d4:5b, Host ID: 80a6d45b.
```

```
ok probe-ide
Device 0 ( Primary Master )
    ATA Model: ST34321A

Device 1 ( Primary Slave )
    Not Present

Device 2 ( Secondary Master )
    Removable ATAPI Model: CRD-8322B

Device 3 ( Secondary Slave )
    Not Present
```

```
ok setenv auto-boot? true
auto-boot? = true
```

Alternatively, you can use the `devalias` command to identify the device aliases and the associated paths of devices that *might* be connected to the system. For example:

```
ok devalias
screen          /pci@1f,0/pci@1,1/SUNW,m64B@2
net             /pci@1f,0/pci@1,1/network@1,1
cdrom          /pci@1f,0/pci@1,1/ide@3/cdrom@2,0:f
disk           /pci@1f,0/pci@1,1/ide@3/disk@0,0
disk3          /pci@1f,0/pci@1,1/ide@3/disk@3,0
disk2          /pci@1f,0/pci@1,1/ide@3/disk@2,0
disk1          /pci@1f,0/pci@1,1/ide@3/disk@1,0
disk0          /pci@1f,0/pci@1,1/ide@3/disk@0,0
ide            /pci@1f,0/pci@1,1/ide@3
floppy         /pci@1f,0/pci@1,1/ibus@1/fdthree
ttyb           /pci@1f,0/pci@1,1/ibus@1/se:b
ttya           /pci@1f,0/pci@1,1/ibus@1/se:a
keyboard!     /pci@1f,0/pci@1,1/ibus@1/su@14,3083f8:forcemode
```

```

keyboard      /pci@1f,0/pci@1,1/ebus@1/su@14,3083f8
mouse         /pci@1f,0/pci@1,1/ebus@1/su@14,3062f8
name          aliases

```

▼ SPARC: How to Determine the Default Boot Device

1 Bring the system to the ok PROM prompt.

For more information, see [“How to Shut Down a Stand-Alone System”](#) on page 199.

2 Use the `printenv` command to determine the default boot device.

```
ok printenv boot-device
```

`boot-device` Identifies the parameter for setting the device from which to boot.

`device[n]` Identifies the `boot-device` value such as a disk or the network. The *n* can be specified as the *disk number*.

The default boot-device is displayed in a format that is similar to the following:

```
boot-device = /pci@1f,4000/scsi@3/disk@1,0:a
```

If the default boot-device is a network boot device, the output is similar to the following:

```
boot-device = /sbus@1f,0/SUNW,fas@e,8800000/sd@a,0:a \
/sbus@1f,0/SUNW,fas@e,8800000/sd@0,0:a disk net
```

▼ SPARC: How to Change the Default Boot Device by Using the Boot PROM

You might need to identify the devices on the system before you can change the default boot device to some other device. For information on identifying devices on the system, see [“SPARC: How to Identify Devices on a System”](#) on page 205.

1 Change to run level 0.

```
# init 0
```

The ok PROM prompt is displayed. For more information, see the `init(1M)` man page.

2 Change the value of the `boot-device` parameter.

```
ok setenv boot-device device[n]
```

Use one of the probe commands if you need help identifying the disk number.

3 Verify that the default boot device has been changed.

```
ok printenv boot-device
```

4 Save the new boot-device value.

```
ok reset-all
```

The new boot-device value is written to the PROM.

Example 11-2 SPARC: Changing the Default Boot Device

In this example, the default boot device is set to disk.

```
# init 0
#
INIT: New run level: 0
.
.
.
The system is down.
syncing file systems... done
Program terminated
ok setenv boot-device /pci@1f,4000/scsi@3/disk@1,0
boot-device = /pci@1f,4000/scsi@3/disk@1,0
ok printenv boot-device
boot-device /pci@1f,4000/scsi@3/disk@1,0
ok boot
Resetting ...

screen not found.
Can't open input device.
Keyboard not present. Using ttya for input and output.

Sun Enterprise 220R (2 X UltraSPARC-II 450MHz), No Keyboard
OpenBoot 3.23, 1024 MB memory installed, Serial #13116682.
Ethernet address 8:0:20:c8:25:a, Host ID: 80c8250a.

Rebooting with command: boot disk1
Boot device: /pci@1f,4000/scsi@3/disk@1,0 File and args:
```

In this example, the default boot device is set to the network.

```
# init 0
#
INIT: New run level: 0
.
.
.
The system is down.
syncing file systems... done
Program terminated
ok setenv boot-device net
boot-device = net
ok printenv boot-device
boot-device net disk
ok reset
Sun Ultra 5/10 UPA/PCI (UltraSPARC-IIi 333MHz), No Keyboard
OpenBoot 3.15, 128 MB memory installed, Serial #number.
Ethernet address number, Host ID: number.
```



```

Boot device: net File and args:
.
.
.
pluto console login:

```

▼ SPARC: How to Change the Default Boot Device by Using the eeprom Command

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Specify the alternate kernel to boot.

```
# eeprom boot-device new-boot-device
```

3 Verify that the new parameter has been set.

```
# eeprom boot-device
```

The output should display the new eeprom value for the boot - device parameter.

SPARC: Resetting the System

Run the following command from the ok prompt:

```
ok reset-all
```

The self-test program, which runs diagnostic tests on the hardware, is executed. Then, if the auto-boot? parameter is set to true, the system is rebooted.

▼ SPARC: How to Change the Default Kernel by Using the Boot PROM

1 Change to run level 0.

```
# init 0
```

The ok PROM prompt is displayed. For more information, see the `init(1M)` man page.

- 2 **Set the `boot-file` property to an alternate kernel.**

```
ok setenv boot-file boot-file
```

- 3 **Verify that the default boot device has been changed.**

```
ok printenv boot-file
```

- 4 **Save the new `boot-file` value.**

```
ok reset-all
```

The new `boot-file` value is written to the PROM.

▼ **SPARC: How to Change the Default Kernel by Using the `eeprom` Command**

- 1 **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

- 2 **Specify the alternate kernel to boot.**

```
# eeprom boot-file new boot-file
```

For example:

```
# eeprom boot-file=kernel.name/sparcv9/unix
```

- 3 **Verify that the new parameter has been set.**

```
# eeprom boot-file
```

The output should display the new `eeprom` value for the specified parameter.

SPARC Support for ITU Construction Tools

The ITU utility is used to convert software packages and patches into driver update (DU) format, and then patch the Oracle Solaris installation media, which enables you to boot your system with these new packages and patches. Starting with Oracle Solaris 10 9/10, the `itu` command has been modified to enable you to create and store an ITU on your SPARC based system, and then boot the system with the newly created ITU.

▼ **SPARC: How to Boot a System With a Newly Created ITU**

The ITU utility enables you to install out-of-box drivers and then apply bug fixes during the installation process. The following procedure describes the process of creating and storing an ITU on a SPARC based system, and then booting the system with the newly created ITU.

1 On the system that you want to boot with an ITU, become superuser.

2 Create the ITU by using the `itu` command with the following options:

```
# itu makedu -r solaris-release -o iso-file package1-name package2-name...
```

<code>makedu</code>	Takes one or more driver packages as input and converts the package to DU format.
<code>-r solaris-release</code>	Specifies the Oracle Solaris release number for which the driver update is intended. The release number format that is used is 5.10.
<code>-o iso-file</code>	Specifies the path of the ISO image file to create.

For more detailed information about the `itu` command and all of its options, see the [itu\(1M\)](#) man page.

3 Bring the system to the `ok PROM` prompt.

4 Disable the `auto-boot` property.

```
ok setenv auto-boot?=false
```

5 Reset the system.

```
ok reset-all
```

Note – The previous step is important, as it ensures that the `itu-device` string is not cleared when you boot the system.

6 Locate the `cdrom` alias.

```
ok devalias cdrom
```

The alias is the device path for the CD or DVD drive that is attached to the system. Note that this output varies, depending on the system. For example, in this procedure, the location of the `cdrom` alias is as follows:

```
/pci@8,700000/scsi@1/disk@6,0:f
```

7 Configure the `itu-device` path that will be used for the installation of the ITU.

Note that the output from Step 6 is used for the first portion of the `config-string`, as shown in the following example:

```
ok> " /pci@8,700000/scsi@1/disk@6,0:f" d# 80 config-string itu-device
```

Note – The 80 refers to the maximum allowable length for the string. If your modified `cdrom` alias is more than 80 characters, increase this number, as appropriate.

8 Boot the system with the ITU.

```
ok boot
```

When you boot the system and the Solaris installation process begins, the system locates the ITU that is stored on the `itu-device`. The system components that are available on the `itu-device` are then identified and configured in the system. The devices for those drivers are also identified and configured. When the installation has completed, the devices become available as installation targets.

9 After the installation has completed, restore the default value for the `auto-boot` property, then boot the system.

```
ok setenv auto-boot?=true
```

The preceding example assumes the default value for the `auto-boot` property is `true`.

10 Boot the system.

```
ok boot
```

Modifying Boot Behavior on x86 Based Systems (Task Map)

Task	Description	For Instructions
Set boot file parameters by using the <code>eeprom</code> command.	Modify boot behavior on an x86 based system by using the <code>eeprom</code> command. Boot options that are set by using the <code>eeprom</code> command persist over a system reboot, unless these options are overridden by modifying kernel behavior in the GRUB menu at boot time.	“x86: How to Modify Boot Behavior by Using the <code>eeprom</code> Command” on page 214

Task	Description	For Instructions
Modify boot behavior by editing the GRUB menu at boot time.	Modify boot behavior by editing GRUB menu at boot time. Boot options that are specified by modifying the boot behavior in the GRUB menu persist only until the next system reboot.	“x86: How to Modify Boot Behavior by Editing the GRUB Menu at Boot Time” on page 217
Modify boot behavior by manually editing the <code>menu.lst</code> file.	Modify boot behavior by editing the <code>menu.lst</code> configuration file to add new OS entries or redirect the console. Changes you make to the file persist over system reboots.	“x86: How to Modify Boot Behavior by Editing the <code>menu.lst</code> File” on page 219
Modify the <code>menu.lst</code> file to include entries that support the <code>findroot</code> command.	Additional menu entries that use the <code>findroot</code> command can be added to the <code>menu.lst</code> file menu after an installation or upgrade.	“x86: How to Add GRUB Menu Entries That Use the <code>findroot</code> Command” on page 224

Modifying Boot Behavior on x86 Based Systems

The primary methods for modifying boot behavior on an x86 based system are as follows:

- By using the `eeprom` command.

The `eeprom` command is used to assign a different value to a standard set of properties. These values, which are equivalent to the SPARC OpenBoot PROM NVRAM variables, are stored in the `/boot/solaris/bootenv.rc` file. Changes that are made to boot behavior by using the `eeprom` command persist over each system reboot and are preserved during a software upgrade. You can override these changes by editing the GRUB menu at boot time or by editing the `menu.lst` file. See the [`eeprom\(1M\)`](#) man page for more information.

Note – Changes that are made by directly editing the `bootenv.rc` file are not always preserved during a software upgrade. This method is therefore discouraged. The preferred method for making these types of changes is to use the `eeprom` command.

- By editing the GRUB menu at boot time.
Changes that are made by modifying the GRUB kernel behavior at boot time override options that you set by using the `eeprom` command. However, these changes only remain in effect until the next time you boot the system. See the [`kernel\(1M\)`](#) man page for more information.
- By manually editing the GRUB `menu.lst` file.



Caution – Any system-generated changes that are made to `menu.lst` entries are changed or lost during a system upgrade. However, any new boot entries that were manually added remain after an upgrade. You can override eeprom settings by editing the GRUB menu at boot time or by editing the `menu.lst` file. Changes made by editing the GRUB menu at boot time do not persist. Whereas, changes that are made by editing the `menu.lst` file persist over system reboots.

▼ x86: How to Modify Boot Behavior by Using the eeprom Command

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Change the specified parameter.

```
# eeprom parameter=new-value
```

3 Verify that the new parameter has been set.

```
# eeprom parameter
```

The output should display the new eeprom value for the specified parameter.

Example 11–3 x86: Setting boot-file Parameters by Using the eeprom Command

This example shows how to manually specify that the system boot a 64-bit kernel. The system must support 64-bit computing.

```
# eeprom boot-file=kernel/amd64/unix
```

This example shows how to manually boot a 32-bit kernel on a 64-bit capable system.

```
# eeprom boot-file=kernel/unix
```

This example shows how to restore the default auto detected boot behavior on a system.

```
# eeprom boot-file=""
```

x86: Modifying Boot Behavior by Editing the GRUB Menu at Boot Time

The following is an example of a GRUB main menu in an Oracle Solaris release that supports booting a system from a ZFS root file system. This menu is based on the contents of the `menu.lst` file, which includes menu entries for all of the bootable OS instances on the system.

The first entry in the menu is the default, unless otherwise specified. To specify another boot entry as the default, add the `default=n` command to the menu. `lst` file, where *n* is a number, starting from 0 (the first boot entry).

```
GNU GRUB  version 0.95  (637K lower / 3144640K upper memory)
+-----+
be1)
be1 failsafe
be3
be3 failsafe
be2
be2 failsafe
+-----+
      Use the ^ and v keys to select which entry is highlighted.
      Press enter to boot the selected OS, 'e' to edit the
      commands before booting, or 'c' for a command-line.
```

Note – The information that is contained in the menu. `lst` file varies, depending on the Oracle Solaris release and the installation method that was used.

To edit a boot entry in the GRUB menu, use the arrow keys to select the entry, then type `e`.

```
GNU GRUB  version 0.95  (637K lower / 3144640K upper memory)
+-----+
findroot (BE_be1,0,a)
bootfs rpool/ROOT/szboot_0508
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive
+-----+
      Use the ^ and v keys to select which entry is highlighted.
      Press enter to boot the selected OS, 'e' to edit the
      commands before booting, or 'c' for a command-line.
```

For instructions on editing the GRUB menu at boot time, see [“x86: How to Modify Boot Behavior by Editing the GRUB Menu at Boot Time” on page 217](#).

Editing the GRUB Menu in Oracle Solaris 10

The following examples show the edit menu in the various GRUB implementations:

GRUB ZFS Support:

```
grub edit> kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS,prop=value
[.prop=value...][-asrvxk] [-m smf-options] [-i altinit]
```

Note – When adding boot arguments on a system with ZFS support, include any additional `-B` options after the default `-B $ZFS -BOOTFS` argument.

GRUB UFS Support:

```
grub edit> kernel /platform/i86pc/multiboot [-asrvxk] [-m smf-options]
[-i altinit][-B prop=value [,prop=value...]]
```

Boot Arguments That Can Be Specified by Editing the GRUB Menu at Boot Time

The following list describes the boot arguments and options that you can specify by editing the GRUB menu at boot time:

<code>multiboot</code>	Specifies the kernel to boot
<code>-a</code>	Prompts the user for configuration information
<code>-s</code>	Boots the system in single-user mode
<code>-r</code>	Specifies a reconfiguration boot
	The system probes all attached hardware devices and then assigns nodes in the file system to represent only those devices that are actually found.
<code>-v</code>	Boots the system with verbose messages enabled
<code>-x</code>	Does not boot the system in clustered mode
<code>-k</code>	Boots the system with the kernel debugger enabled
<code>-m smf-options</code>	Controls the boot behavior of the Service Management Facility (SMF)
	There are two categories of options: recovery options and messages options.
<code>-i altinit</code>	Specifies an alternative executable as the primordial process. <code>altinit</code> is a valid path to an executable.
<code>-B prop=value [,prop=value]...</code>	Specifies kernel boot properties.

The following are various ways you can modify boot behavior in the GRUB menu by using the `-B prop=val` option:

`-B console=ttya` Redirects the console to `ttya`.

-B acpi-enum=off	Disables Advanced Configuration and Power Interface (ACPI) enumeration of devices.
-B console=ttya,acpi-enum=off	Redirects the console to ttya and disables the ACPI enumeration of devices.
-B acpi-user-options=0x2	Disables ACPI entirely.

Note – When properties are specified by using the `eeprom` command *and* on the GRUB command line, the GRUB command takes precedence.

▼ x86: How to Modify Boot Behavior by Editing the GRUB Menu at Boot Time

When you modify the GRUB kernel behavior by editing the GRUB menu at boot time, the changes do not persist over a system reboot. Default boot behavior is restored the next time you boot the system.

- 1 **Reboot the system.**
When the boot sequence begins, the GRUB main menu is displayed.
- 2 **Use the arrow keys to select the boot entry to edit.**
- 3 **Type `e` to access the GRUB edit menu.**
- 4 **Select the `kernel` or `kernel$` line in this menu.**
- 5 **Type `e` to add boot arguments to the line.**
- 6 **Type any additional boot arguments.**
- 7 **Press Return to save your changes and return to the previous menu.**

Note – Pressing the Escape key returns you to the GRUB main menu without saving your changes.

- 8 **To boot the system, type `b`.**
Changes you make take affect when the system is booted.

Example 11-4 x86: Booting a 32-Bit Kernel on a 64-Bit Enabled System

To boot a 32-bit kernel on a 64-bit capable system, add the `kernel/unix` argument.

```
grub edit> kernel /platform/i86pc/multiboot kernel/unix
```

Example 11-5 x86: Redirecting the Serial Console

To redirect the serial console to `ttyb`, add the `-B console=ttyb` argument.

```
grub edit> kernel /platform/i86pc/multiboot -B console=ttyb
```

Alternatively, you can use `input-device/output-device` property, as shown in the following example:

```
grub edit> kernel /platform/i86pc/multiboot -B input-device=ttyb,output-device=ttyb
```

This example shows how you would override the serial line speed:

```
grub edit> kernel /platform/i86pc/multiboot -B ttyb-mode="115200,8,n,1,-"
```

Caution: In the preceding example, the property value contains commas, which is also a property separator. To avoid confusing the property parser, use double quotation marks around the entire property value.

x86: Modifying Boot Behavior by Editing the `menu.lst` File

The GRUB menu, which is based on the `menu.lst` configuration file, can be customized. The operating systems that are installed on your system are listed in this file and displayed in the GRUB menu when the system is booted. Note that when you install an operating system other than Oracle Solaris, you must manually add a menu entry for that operating system to the `menu.lst` file after installation.

The following is an example of a typical GRUB main menu, which is based on the contents of the `menu.lst` file. The GRUB main menu consists of all of the boot entries that are available, plus a failsafe archive.

```
GNU GRUB version 0.95 (631K lower / 2095488K upper memory)
+-----+
| Solaris 10.1 ... X86                               |
| Solaris failsafe                                   |
|                                                    |
+-----+
```

A configurable timeout is available to boot the default OS entry. The default OS boot entry that is booted is configurable through the `default` command. The installation software typically sets this command to boot one of the valid boot entries. To boot a different instance of the Oracle Solaris OS (if applicable), or to boot a different OS, use the arrow keys to highlight a different boot entry. Then press Enter to boot that entry. Note that if the `default` command is not set, the first boot entry in the GRUB menu is booted.

Only the *active* `menu.lst` file is used to boot the system. To modify the GRUB menu that is displayed when you boot the system, edit the active GRUB `menu.lst` file. Changing any other `menu.lst` file has no effect on the menu that is displayed when you boot the system. To determine the location of the active `menu.lst` file, use the `list-menu` subcommand of the `bootadm` command. For more information about using the `bootadm` command, see [“Using the bootadm Command to Manage the Boot Archives” on page 272](#).

For a complete description of the `menu.lst` file in each of the GRUB implementations in the Solaris OS, see [“x86: Supported GRUB Versions” on page 295](#).

▼ x86: How to Modify Boot Behavior by Editing the menu.lst File

You might need to modify the `menu.lst` file for one of the following reasons:

- To add new OS entries
- To add GRUB console redirection information

Before You Begin Because only the *active* GRUB `menu.lst` file is used to boot the system, make sure you edit the correct file. Changing any other GRUB `menu.lst` file has no effect on the menu that is displayed when you boot the system.

The location of the active `menu.lst` file varies, depending on whether you have a system with a UFS root or a ZFS root.

- For a UFS root, the active `menu.lst` file is `/boot/grub/menu.lst`.
- For a ZFS root, the active `menu.lst` file is `/pool-name/boot/grub/menu.lst`.

You can determine the location of the active GRUB `menu.lst` file by using the `bootadm` command with the `list-menu` subcommand.

```
# bootadm list-menu
```

For more information about the `bootadm` command, see the [`bootadm\(1M\)`](#) man page.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 To add a new OS entry to the active menu .lst file, use a text editor to modify the file.

The comments within the menu .lst file provide you with the necessary information for adding a new OS entry.

The following is an example of a menu .lst file for a system that is running a release with ZFS boot support. Boot entries in the menu .lst file vary, depending on the Oracle Solaris release that you are running.

```
#----- ADDED BY BOOTADM - DO NOT EDIT -----  
title Solaris Solaris 10 s10x_nbu6wos_nightly X86  
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS  
module /platform/i86pc/boot_archive  
#-----END BOOTADM-----
```



Caution – Do not directly edit the original contents of the menu .lst file. To make changes to any of the OS entries in the file, manually edit the file to duplicate the existing content. Then, make the modifications to the duplicated content.

Also note when manually adding new entries to the file, never include guard comments, for example, “Added by bootadm”. These comments are reserved for use by the system. Not using these comments ensures that these entries remain intact during a software upgrade.

If you have added any additional entries, beyond the default entries, make any equivalent changes manually.

The [-B*] and [*] flags must be preserved, if these flags exist in the original menu .lst file. Also, the failsafe entry should always have an -s flag.

3 After adding the required information, save the file.

Note that any changes you make to the file take effect at the next system reboot.

Tip – If you are running Linux, and install Oracle Solaris, the Linux entry is not preserved in the GRUB menu when the system is rebooted. Before installing or upgrading your system, save a copy of the menu .lst file that contains the Linux information. After the installation, add the Linux information to the newly created menu .lst file in the Solaris partition.

Because changes you make to the menu .lst file are not directly related to the Oracle Solaris OS, you cannot make them by using the eeprom command. You must edit the file directly. Note that the software upgrade process preserves any changes that you make to the menu .lst file.



Caution – GRUB is capable of booting both Linux and Oracle Solaris. However, Linux GRUB is not capable of booting Oracle Solaris.

Always ensure that one of the following conditions are met:

- That the `fdisk` partition is active, that it has GRUB installed, and that the `menu.lst` file is the *active* GRUB menu.
- That Oracle Solaris GRUB is installed to the Master Boot Record (MBR), and that it refers to the `menu.lst` in the `fdisk` partition.

For a detailed description of the GRUB `menu.lst` that pertains to each Oracle Solaris release, see [“x86: Supported GRUB Versions” on page 295](#).

Example 11–6 `menu.lst` File on a System With an Oracle Solaris ZFS Boot Loader

The following examples show what a `menu.lst` file looks like on a system that has an Oracle Solaris ZFS boot loader. By default, this system will boot from a ZFS root file system. Note that the contents of the file varies, depending on the installation type.

New installation or standard upgrade:

```
title Solaris 10 s10x_nbu6wos_nightly X86
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title Solaris failsafe
findroot (pool_rpool,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```

Oracle Solaris Live Upgrade

```
title bel
findroot (BE_bel,0,a)
bootfs rpool/ROOT/szboot_0508
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title bel failsafe
findroot (BE_bel,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-sa
```

Example 11–7 `menu.lst` File on a System With a UFS Boot Loader

The following examples show what a `menu.lst` file looks like on a system that has a UFS root file system installed. By default, this system will boot from a UFS root file system.

New installation or standard upgrade:

```
title Solaris 10 s10x_nbu6wos_nightly X86
findroot (rootfs0,0,a)
kernel /platform/i86pc/multiboot
module /platform/i86pc/boot_archive

title Solaris failsafe
findroot (rootfs0,0,a)
kernel /boot/multiboot kernel/unix -s      -B console=ttyb
module /boot/x86.miniroot-safe
```

Oracle Solaris Live Upgrade:

```
title be1
findroot (BE_be1,0,a)
kernel /platform/i86pc/multiboot
module /platform/i86pc/boot_archive

title be1 failsafe
findroot (BE_be1,0,a)
kernel /boot/multiboot kernel/unix -s      -B console=ttyb
module /boot/x86.miniroot-safe
```

x86: Locating the Active GRUB menu .lst File

On systems that have a ZFS root, the active menu .lst file is typically located in */pool-name/boot/grub/menu.lst*.

On systems that have a UFS root, the active menu .lst file is typically located in */boot/grub/menu.lst*.

To locate the active GRUB menu, use the `bootadm` command with the `list -menu` subcommand:

```
# bootadm list-menu
```

This command also lists the contents of the active menu .lst file:

```
# bootadm list-menu
The location for the active GRUB menu is: /pool-name/boot/grub/menu.lst
default 0
timeout 10
0 be1
1 be1 failsafe
2 be3
3 be3 failsafe
4 be2
5 be2 failsafe
```

For further instructions on using the `bootadm` command, see [“Using the bootadm Command to Manage the Boot Archives”](#) on page 272.

x86: Implementation of the findroot Command

All installation methods, including Oracle Solaris Live Upgrade, now use the `findroot` command for specifying which disk slice on an x86 based system to boot. This enhancement supports booting systems with Oracle Solaris ZFS roots, as well as UFS roots. This information is located in the `menu.lst` file that is used by GRUB. Previously, the `root` command, `root (hd0.0.a)`, was explicitly used to specify which disk slice to boot.

The installation methods include Oracle Solaris Live Upgrade, Oracle Solaris JumpStart, and the installation GUI program.

In addition to the `findroot` command, is a signature file on the slice, (`mysign, 0, a`), where `mysign` is the name of a signature file that is located in the `/boot/grub/bootsign` directory. When booting a system from a ZFS root, the ZFS GRUB plug-in looks for and tries to mount a ZFS file system in slice `a` of `fdisk` partition `0`.

The name of the signature file varies, depending on the installation method that was used. For more information about the naming conventions that are used by the `findroot` command, see [“Naming Conventions That Are Used by the findroot Command” on page 293](#).

Additional menu entries that also use the `findroot` command can be added to the GRUB menu after an installation or upgrade. For instructions, see [“x86: How to Add GRUB Menu Entries That Use the findroot Command” on page 224](#).



Caution – The boot signature must be unique. Do not use or remove system-generated signatures or user signatures that are duplicated across multiple instances of the Oracle Solaris software. Doing so might result in booting an incorrect OS instance or prevent the system from booting.

Note that the `root` command can still be used in the `menu.lst` file in certain instances, for example to boot Windows. However, do not use the `root` command in cases where the `findroot` command is the preferred choice.

EXAMPLE 11-8 x86: Default `menu.lst` file on a System That Supports a UFS Boot Loader

The following example shows the format of a `menu.lst` file entry that uses the `findroot` command:

```
title Solaris 10 s10x_nbu6wos_nightly X86
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title Solaris failsafe
findroot (pool_rpool,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```

EXAMPLE 11-9 x86: Default menu.lst file That Supports an Oracle Solaris ZFS Boot Loader

This is an example of a menu.lst file on system that supports an Oracle Solaris ZFS boot loader. The information for booting from a ZFS root file system is automatically added to the file when an Oracle Solaris Live Upgrade is performed.

```
title bel
findroot (BE_bel,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title bel failsafe
findroot (BE_bel,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```

▼ x86: How to Add GRUB Menu Entries That Use the findroot Command

This procedure shows how to manually update the menu.lst file with user-defined entries that use the findroot command. Typically, these entries are added after an installation or an upgrade. For guidelines on adding user-defined entries that use the findroot command, see “x86: Implementation of the findroot Command” on page 223.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Create a boot signature file on the root (/) file system or root pool that will be booted.

- For a ZFS pool, *my-pool*, create the boot signature file in the */my-pool/boot/grub/bootsign* directory.

```
# touch /my-pool/boot/grub/bootsign/user-sign
```

- For a UFS file system, create the boot signature file in the */boot/grub/bootsign* directory of the root file system to be booted.

```
# touch /boot/grub/bootsign/user-sign
```

Note – Make sure the file name that you choose for the boot signature is unique. Do not use system-generated signature names or user signature names that are duplicated across multiple instances of Oracle Solaris. Doing so might prevent the system from booting or cause the wrong Oracle Solaris instance to boot.

3 Add a menu entry that contains the `findroot` command.**a. Locate the active menu .lst file:**

```
# bootadm list-menu
```

b. Using a text editor, edit the active menu .lst file to add the following entry:

```
title    User Solaris boot entry
findroot (user-sign, 3, c)
kernel$  /platform/i86pc/multiboot
module   /platform/i86pc/boot_archive
```

In the preceding example, the 3 represents the 4th fdisk partition (partitions start at 0). The c represents the slice within a Solaris fdisk partition (slices start with a).

4 Reboot the system.

The new entry is displayed in the GRUB menu and can be selected to boot the specified Oracle Solaris OS instance.

Booting an Oracle Solaris System (Tasks)

This chapter describes the procedures for booting the Oracle Solaris release on SPARC and x86 based systems.

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

- “Booting a SPARC Based System (Task Map)” on page 228
- “Booting a SPARC Based System” on page 229
- “Booting From a Specified ZFS Root File System on a SPARC Based System” on page 234
- “Booting a SPARC Based System in Failsafe Mode” on page 238
- “Booting a SPARC Based System From the Network” on page 243
- “Booting an x86 Based System by Using GRUB (Task Map)” on page 245
- “x86: Booting From a Specified ZFS Root File System on an x86 Based System” on page 252
- “Booting an x86 Based System in Failsafe Mode” on page 255
- “Booting an x86 Based System From the Network” on page 260
- “Accelerating the Reboot Process on the SPARC Platform (Task Map)” on page 264
- “Initiating a Fast Reboot of a SPARC Based System” on page 265

For overview information about the boot process, see [Chapter 9, “Shutting Down and Booting a System \(Overview\)”](#).

Note – Starting with the Solaris 10 1/06 release, the open source GRand Unified Bootloader (GRUB) has been implemented on x86 based systems. GRUB is responsible for loading a boot archive, which contains the kernel modules and configuration files, into the system's memory.

For information about booting an x86 based system in a Solaris release that does *not* implement GRUB based booting, see [Chapter 16, “x86: Booting a System That Does Not Implement GRUB \(Tasks\)”](#).

Booting a SPARC Based System (Task Map)

Task	Description	For Instructions
Boot a SPARC based system to run level 3.	Use this boot method after shutting down the system or performing a system hardware maintenance task.	“SPARC: How to Boot a System to Run Level 3 (Multiuser Level)” on page 229
Boot a SPARC based system to run level S.	Use this boot method to boot the system after performing a system maintenance task such as backing up a file system. At this level, only local file systems are mounted and users cannot log in to the system.	“SPARC: How to Boot a System to Run Level S (Single-User Level)” on page 230
Boot a SPARC based system interactively.	Use this boot method after making temporary changes to a system file or the kernel for testing purposes.	“SPARC: How to Boot a System Interactively” on page 231
Boot a Solaris kernel other than default.	Use this procedure to boot a Solaris kernel other than the default kernel. Alternately, you can obtain a copy of an alternate boot file, change the default kernel to the new kernel, then set the <code>boot - file</code> parameter to boot the new default boot device.	“SPARC: How to Boot a Kernel Other Than the Default Kernel” on page 232
Display a list of the available ZFS bootable datasets on a SPARC based system.	Use the <code>boot - L</code> command to display a list of the available BEs within a ZFS pool on a system. Note – This option is only supported for boot devices that contain a ZFS pool.	“SPARC: How to List Available Bootable Datasets Within a ZFS Root Pool” on page 235
Boot a SPARC based system from a ZFS root file system.	Use the <code>boot - Z</code> option to boot a specified ZFS dataset. Note – This option is only supported for boot devices that contain a ZFS pool.	“SPARC: How to Boot From a Specified ZFS Root File System” on page 236
Boot the failsafe archive on a SPARC based system.	Use this procedure to boot a SPARC based system in failsafe mode. Then, run the <code>bootadm</code> command to update the boot archive.	“How to Boot a SPARC Based System in Failsafe Mode” on page 239
Boot a SPARC based system from the network.	Use this boot method to boot a system from the network. Note that this method is also used for booting a diskless client.	“SPARC: How to Boot a System From the Network” on page 243

Booting a SPARC Based System

If a system is turned off, turning it on starts the multiuser boot sequence. The following procedures show how to boot to different run levels from the ok PROM prompt. These procedures assume that the system has been cleanly shut down, unless stated otherwise.

Use the `who -r` command to verify that the system is brought to the specified run level. For a description of run levels, see [Chapter 18, “Managing Services \(Overview\)”](#).

▼ SPARC: How to Boot a System to Run Level 3 (Multiuser Level)

Use this procedure to boot a system that is currently at run level 0 to run level 3.

1 Boot the system to run level 3.

```
ok boot
```

The automatic boot procedure displays a series of startup messages, and brings the system to run level 3. For more information, see the [boot\(1M\)](#) man page.

2 Verify that the system has booted to run level 3.

The login prompt is displayed when the boot process has finished successfully.

```
hostname console login:
```

Example 12–1 SPARC: Booting a System to Run Level 3 (Multiuser Level)

The following example displays the messages from booting a system to run level 3.

```
ok boot
Sun Ultra 5/10 UPA/PCI (UltraSPARC-III 333MHz)
OpenBoot 3.15, 128 MB memory installed, Serial #number.
Ethernet address number, Host ID: number.

Rebooting with command: boot
Boot device: /pci@1f,0/pci@1,1/ide@3/disk@0,0:a File and args: kernel/sparcv9/unix
SunOS Release 5.10 Version s10_60 64-bit
Copyright 1983-2004 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
configuring IPv4 interfaces: hme0.
add net default: gateway 172.20.27.248
Hostname: starlite
The system is coming up. Please wait.
NIS domain name is example.com
starting rpc services: rpcbind keyser vpbind done.
Setting netmask of hme0 to 255.255.255.0
Setting default IPv4 interface for multicast: add net 224.0/4: gateway starlite
syslog service starting.The system is ready.
```

```
Starting Sun(TM) Web Console Version 2.1-dev..
volume management starting.
The system is ready.
starlite console login:
```

In the preceding example, *sparcv9* was used as an example only. This string matches the output of the `isainfo -k` command.

▼ SPARC: How to Boot a System to Run Level S (Single-User Level)

Use this procedure to boot a system that is currently at run level 0 to run level S. This run level is used for system maintenance tasks, such as backing up a file system.

1 Boot the system to run level S.

```
ok boot -s
```

2 Type the superuser password when the following message is displayed:

```
SINGLE USER MODE
```

```
Root password for system maintenance (control-d to bypass): xxxxxx
```

3 Verify that the system is at run level S.

```
# who -r
```

4 Perform the maintenance task that required the run level change to S.

5 After you complete the system maintenance task, type Control-D to bring the system to the multiuser state.

Example 12–2 SPARC: Booting a System to Run Level S (Single-User Level)

The following example displays the messages from booting a system to run level S.

```
ok boot -s
.
.
.
Sun Microsystems Inc. SunOS 5.10 Version Generic_120012-14 32-bit
Copyright 1983-2003 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
configuring IPv4 interfaces: hme0.
Hostname: starlite

SINGLE USER MODE

Root password for system maintenance (control-d to bypass): xxxxxx
```

```

single-user privilege assigned to /dev/console.
Entering System Maintenance Mode
Oct 14 15:01:28 su: 'su root' succeeded for root on /dev/console
Sun Microsystems Inc. SunOS 5.10
# who -r
. run-level S Sep 19 08:49 S 0 ?
(Perform some maintenance task)
# ^D

```

▼ SPARC: How to Boot a System Interactively

Use this boot option when you need to specify an alternate kernel or `/etc/system` file.

Before You Begin To specify an alternate `/etc/system` file when booting a SPARC based system interactively by using the `boot -a` command, you must perform the following steps before the system is booted.

- 1. Make backup copies of the `/etc/system` and `boot/solaris/filelist.ramdisk` files.

```
# cp /etc/system /etc/system.bak
# cp /boot/solaris/filelist.ramdisk /boot/solaris/filelist.ramdisk.orig
```

- 2. Add the `etc/system.bak` file name to the `/boot/solaris/filelist.ramdisk` file.

```
# echo "etc/system.bak" >> /boot/solaris/filelist.ramdisk
```

- 3. Update the boot archive.

```
# bootadm update-archive -v
```

1 Boot the system interactively.

```
ok boot -a
```

2 Answer the following system prompts:

a. When prompted, enter the name of the kernel to use for booting.

Press enter to use the default kernel file name. Otherwise, provide the name of an alternate kernel, press Enter.

b. When prompted, provide an alternate path for the `modules` directories.

Press enter to use the default module directories. Otherwise, provide the alternate paths to module directories, press Enter.

c. When prompted, provide the name of an alternate system file.

Type `/dev/null` if your `/etc/system` file has been damaged.

d. When prompted, enter the root filesystem type.

Press enter to select UFS for local disk booting, which is the default, or enter NFS for network booting.

e. **When prompted, enter the physical name of root device.**

Provide an alternate device name or press return to use the default.

3 If you are not prompted to answer these questions, verify that you typed the boot -a command correctly.

Example 12-3 SPARC: Booting a System Interactively

In this example, the default choices (shown in square brackets []) are accepted. For instructions and an example of booting an alternate file system by using the boot -a command, see “[SPARC: How to Boot a System Interactively](#)” on page 231.

```
ok boot -a
.
.
.
Rebooting with command: boot -a
Boot device: /pci@1f,0/pci@1,1/ide@3/disk@0,0:a
File and args: -a
Enter filename [kernel/sparcv9/unix]:      Press Return
Enter default directory for modules [/platform/SUNW,Ultra-5_10/kernel
/platform/sun4u/kernel /kernel /usr/kernel]:      Press Return
Name of system file [etc/system]:      Press Return
SunOS Release 5.10 Version S10_60 64-bit
Copyright (c) 1983-2004 by Sun Microsystems, Inc. All rights reserved
Use is subject to license terms.
root filesystem type [ufs]:      Press Return
Enter physical name of root device
[/pci@1f,0/pci@1,1/ide@3/disk@0,0:a]:      Press Return
configuring IPv4 interfaces: hme0.
Hostname: starlite
The system is coming up. Please wait.
checking ufs filesystems
.
.
.
The system is ready.
starlite console login:
```

▼ SPARC: How to Boot a Kernel Other Than the Default Kernel

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

- 2 Obtain a copy of an existing Oracle Solaris kernel and rename it.
- 3 Add the kernel that you copied and renamed in Step 2 to the `/etc/boot/solaris/filelist.ramdisk` file.


```
# echo "kernel.name" >> /boot/solaris/filelist.ramdisk
```
- 4 Verify that the alternate kernel has been added to the `/etc/boot/solaris/filelist.ramdisk` file.


```
# cat > /etc/boot/solaris/filelist.ramdisk
```
- 5 Update the boot archive by using the `bootadm` command.


```
# bootadm update-archive
```
- 6 Change to run level 0.


```
# init 0
```

The ok PROM prompt is displayed.
- 7 Boot the alternate kernel.


```
ok boot alternate-kernel
```

For example:

```
ok boot kernel.myname/sparcv9/unix
```

 - To boot the alternate kernel by default, follow these steps:
 - a. Set the boot-file parameter to the new kernel.


```
ok setenv boot-file kernel.name/sparcv9/unix
```
 - b. Verify that the boot-file property has been changed.


```
ok printenv boot-file
```
 - c. Reboot the system.


```
ok boot
```
- 8 After the system has booted, verify that the alternate kernel that was booted.


```
# prtconf -vp | grep whoami
```

Example 12-4 Booting an Alternate Kernel by Changing the Default Boot File

```
# cp -r /platform/sun4v/kernel /platform/sun4vu/kernel.caiobella
# echo "kernel.caiobella" >> /boot/solaris/filelist.ramdisk

# cat > /etc/boot/solaris/filelist.ramdisk
/platform/sun4v/kernel.caiobella
^D (control D)
```

```
ok setenv boot-file kernel.caibells/sparcv9/unix
ok printenv boot-file
boot-file = kernel.caibella/sparcv9/unix
```

```
ok boot
```

```
SC Alert: Host System has Reset
```

```
SC Alert: Host system has shut down.
```

```
Sun Fire T200, No KeyboardCopyright 2006 Sun Microsystems, Inc. All rights reserved.
OpenBoot 4.25.0.build_01***PROTOTYPE BUILD***, 32760 MB memory available, Serial
#69060038.
Ethernet address 0:x:4f:x:c5:c6, Host ID: 8xxc5c6.
```

```
Rebooting with command: boot
Boot device: /pci@7c0/pci@0/pci@1/pci@0,2/LSILogic,sas@2/disk@0,0:a File and
args: kernel.caibella/sparcv9/unix
SunOS Release 5.10
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
DEBUG enabled
misc/forthdebug (176650 bytes) loaded
Hostname: seasonz
NIS domain name is lab.domain.sun.com
Reading ZFS config: done.
```

```
seasonz console login:
Password:
Last login: Mon Nov 12 18:02:00 on console
Sun Microsystems Inc. SunOS 5.10
.
.
.
You have new mail.
#
#
# prtconf -vp | grep whoami
whoami: '/platform/sun4v/kernel.caibella/sparcv9/unix'
```

Booting From a Specified ZFS Root File System on a SPARC Based System

To support booting from an Oracle Solaris ZFS on the SPARC platform, two new boot options have been added:

-L Displays a list of available bootable datasets within a ZFS pool.

Note – The `boot -L` command is executed from the OBP, *not* from the command line.

`-Z dataset` Boots the root file system for the specified ZFS bootable dataset.

If you are booting a system from a ZFS root file system, first use the `boot` command with the `-L` option from the OBP to print a list of the available BEs on the system. Then, use the `-Z` option to boot the specified BE.

For more information, see the [boot\(1M\)](#) man page.

▼ SPARC: How to List Available Bootable Datasets Within a ZFS Root Pool

On SPARC based systems, the `menu.lst` file contains the following two GRUB commands:

- `title` – Provides a title for a boot environment (BE)
- `bootfs` – Specifies the full name of the bootable dataset

To display a list of the bootable datasets within a ZFS pool, choose from the following methods:

- Use the `lustatus` command. This command lists all of the BEs in a given ZFS pool. Note that the `lustatus` command can also be used on x86 based systems.
- Use the `boot -L` command. This command displays a list of the available BEs in a given ZFS pool and provides instructions for booting the system.

The following procedure describes how to use the `boot -L` command to list available BEs on a system. To boot a specified BE after running this command, follow the instructions that are printed on the screen.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

2 Bring the system to the `ok PROM` prompt.

```
# init 0
```

3 List the available BEs in a ZFS pool:

```
ok boot device-specifier -L
```

4 To boot one of the entries that is displayed, type the number that corresponds to the entry.

5 Boot the specified BE by following the directions that are printed on the screen.

For instructions, see “[SPARC: How to Boot From a Specified ZFS Root File System](#)” on page 236.

Example 12-5 SPARC: Displaying a List of Available BEs on a System by Using boot -L

```
# init 0
# svc.startd: The system is coming down. Please wait.
svc.startd: 94 system services are now being stopped.
svc.startd: The system is down.
syncing file systems... done
Program terminated
ok boot -L
.
.
.
Boot device: /pci@1f,0/pci@1/scsi@8/disk@0,0 File and args: -L
zfs-file-system
Loading: /platformsun4u/bootlst
1.s10s_nbu6wos
2 zfs2BE
Select environment to boot: [ 1 - 2 ]: 2

to boot the selected entry, invoke:
boot [<root-device>] -Z rpool/ROOT/zfs2BE
```

See Also For more information, see [Chapter 5, “Installing and Booting an Oracle Solaris ZFS Root File System,”](#) in *Oracle Solaris ZFS Administration Guide*.

▼ SPARC: How to Boot From a Specified ZFS Root File System

Booting from Oracle Solaris ZFS differs from booting from UFS. When booting from ZFS, a device specifier identifies a storage pool, *not* a single root file system. A storage pool can contain multiple bootable datasets, or root file systems. Therefore, when booting from ZFS, you must also identify a root file system within the pool that is identified by the boot device as the default. By default, the default boot device is identified by the pool's boot fs property. This procedure shows how to boot the system by specifying a ZFS bootable dataset. See the [boot\(1M\)](#) man page for a complete description of all the boot options that are available.

Note – If the boot fs property was previously set up correctly, for example, if you used the `luactivate` command to activate a BE, the system boots a ZFS root automatically.

For more information, see [zpool\(1M\)](#) man page.

- 1 **Become superuser or assume an equivalent role.**
Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.
- 2 **Bring the system to the ok PROM prompt.**
`# init 0`
- 3 **(Optional) To display a list of available BEs, use the boot command with the -L option.**
For instructions, see [“SPARC: How to List Available Bootable Datasets Within a ZFS Root Pool”](#) on page 235.
- 4 **To boot a specified entry, type the number of the entry and press Return:**
Select environment to boot: [1 - 2]:
- 5 **To boot the system, follow the instructions that are printed to the screen:**
To boot the selected entry, invoke:
`boot [<root-device>] -Z rpool/ROOT/dataset`

`ok boot -Z rpool/ROOT/dataset`
For example:

`# boot -Z rpool/ROOT/zfs2BE`
- 6 **After the system has booted, type the following command to verify the active BE:**
`# prtconf -vp | grep whoami`
 - **To display the boot path for the active BE, type:**
`# prtconf -vp | grep bootpath`
 - **Alternately, you can use the df -lk command to determine whether the correct BE was booted.**

Example 12–6 SPARC: Booting From a Specified ZFS Root File System

This example shows how to use the `boot -Z` command to boot a ZFS dataset on a SPARC based system.

```
# init 0
# svc.startd: The system is coming down. Please wait.
svc.startd: 79 system services are now being stopped.
svc.startd: The system is down.
syncing file systems... done
Program terminated
ok boot -Z rpool/ROOT/zfs2BEe
Resetting
```

```

LOM event: =44d+21h38m12s host reset
g ...

rProcessor Speed = 648 MHz
Baud rate is 9600
8 Data bits, 1 stop bits, no parity (configured from lom)

Firmware CORE Sun Microsystems, Inc.
@(#) core 1.0.12 2002/01/08 13:00
software Power ON
Verifying nVRAM...Done
Bootmode is 0
[New I2C DIMM address]
.
.
.
Environment monitoring: disabled
Executing last command: boot -Z rpool/ROOT/zfs2BE
Boot device: /pci@1f,0/pci@1/scsi@8/disk@0,0 File and args: -Z rpool/ROOT/zfs2Be
zfs-file-system
Loading: /platform/SUNW,UltraAX-i2/boot_archive
Loading: /platform/sun4u/boot_archive
ramdisk-root hsfs-file-system
Loading: /platform/SUNW,UltraAX-i2/kernel/sparcv9/unix
Loading: /platform/sun4u/kernel/sparcv9/unix
.
.
.
Hostname: mallory
NIS domainname is boulder.Central.Sun.COM
Reading ZFS config: done.
Mounting ZFS filesystems: (6/6)

mallory console login:

```

See Also For information about booting the failsafe archive for a specified ZFS bootable dataset, see [“How to Boot a SPARC Based System in Failsafe Mode” on page 239](#).

Booting a SPARC Based System in Failsafe Mode

Booting a system from a root file system image that is a boot archive, and then remounting this file system on the actual root device can sometimes result in a boot archive and root file system that do not match, or are *inconsistent*. Under these conditions, the proper operation and integrity of the system is compromised. After the root (/) file system is mounted, and before relinquishing the in-memory file system, the system performs a consistency verification against the two file systems. If an inconsistency is detected, the normal boot sequence is suspended and the system reverts to *failsafe mode*.

Also, if a system failure, a power failure, or a kernel panic occurs immediately following a kernel file update, the boot archives and the root (file system might not be synchronized. Although the

system might still boot with the inconsistent boot archives, it is recommended that you boot the failsafe archive to update the boot archives. You can also use the `bootadm` command to manually update the boot archives. For more information, see [“Using the bootadm Command to Manage the Boot Archives” on page 272](#).

The failsafe archive can be booted for recovery purposes, or to update the boot archive.

On the SPARC platform the failsafe archive is:

```
/platform/‘uname -m’/failsafe
```

You would boot the failsafe archive by using the following syntax:

```
ok boot -F failsafe
```

Failsafe booting is also supported on systems that are booted from Oracle Solaris ZFS. When booting from a ZFS-rooted BE, each BE has its own failsafe archive. The failsafe archive is located where the root file system is located, as is the case with a UFS-rooted BE. The default failsafe archive is the archive that is in the default bootable file system. The default bootable file system (dataset) is indicated by the value of the pool's `bootfs` property.

For information about booting an x86 based failsafe archive, see [“Booting an x86 Based System in Failsafe Mode” on page 255](#).

For information about clearing automatic boot archive failures, see [“x86: How to Clear Automatic Boot Archive Update Failures by Using the auto-reboot-safe Property” on page 270](#).

▼ How to Boot a SPARC Based System in Failsafe Mode

Use this procedure to boot a SPARC based system in failsafe mode for the purpose of updating the boot archive. If the system does not boot after the boot archive is updated, you might need to boot the system in single-user mode. For more information, see [“SPARC: How to Boot a System to Run Level S \(Single-User Level\)” on page 230](#).

Note – This procedure also provides instructions for booting a specific ZFS dataset in failsafe mode.

1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in *System Administration Guide: Security Services*](#).

2 Bring the system to the ok prompt:

```
# init 0
```

3 Boot the failsafe archive.

- **To boot the default failsafe archive, type:**

```
ok boot -F failsafe
```

- **To boot the failsafe archive of a specific ZFS dataset:**

```
ok boot -F failsafe -Z dataset
```

For example:

```
ok boot -F failsafe -Z rpool/ROOT/zfsBE2
```

Note – To determine the name of the dataset to boot, first use the boot -L command to display a list of the available BEs on the system. For more information, see [“SPARC: How to List Available Bootable Datasets Within a ZFS Root Pool” on page 235.](#)

If an inconsistent boot archive is detected a message is displayed.

4 To update the boot archive, type y, then press Return.

An out of sync boot archive was detected on rpool.
The boot archive is a cache of files used during boot
and should be kept in sync to ensure proper system operation.

Do you wish to automatically update this boot archive? [y,n,?] **y**

If the archive was updated successfully, a message is displayed:

The boot archive on rpool was updated successfully.

Example 12-7 SPARC: Booting a System in Failsafe Mode

This example shows how to boot a SPARC based system in failsafe mode. If no device is specified, the failsafe archive for the default boot device is booted.

```
ok boot -F failsafe
Resetting ...
screen not found.
Can't open input device. Keyboard not present. Using ttya for input and output.
```

```
Sun Enterprise 220R (2 X UltraSPARC-II 450MHz), No Keyboard
OpenBoot 3.23, 1024 MB memory installed, Serial #13116682.
Ethernet address 8:0:20:c8:25:a, Host ID: 80c8250a.
```

```
Rebooting with command: boot -F failsafe
Boot device: /pci@1f,4000/scsi@3/disk@1,0:a File and args: -F failsafe
```



```

SunOS Release 5.10t
Copyright 1983-2007 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Configuring /dev Searching for installed OS instances...

An out of sync boot archive was detected on /dev/dsk/c0t1d0s0.
The boot archive is a cache of files used during boot and
should be kept in syncto ensure proper system operation.

Do you wish to automatically update this boot archive? [y,n,?] y
Updating boot archive on /dev/dsk/c0t1d0s0.
The boot archive on /dev/dsk/c0t1d0s0 was updated successfully.

Solaris 5.10 was found on /dev/dsk/c0t1d0s0.
Do you wish to have it mounted read-write on /a? [y,n,?] n
Starting shell.
#

```

Example 12-8 SPARC: Booting a Specified ZFS Dataset in Failsafe Mode

This example shows how to boot the a ZFS dataset in failsafe mode. Note that the boot -L command is first used to display a list of available boot environments. This command must be run at the ok prompt.

```

ok boot -L
Rebooting with command: boot -L
Boot device: /pci@1f,4000/scsi@3/disk@1,0 File and args: -L
1 zfsBE2
Select environment to boot: [ 1 - 1 ]: 1

To boot the selected entry, invoke:
boot [<root-device>] -Z rpool/ROOT/zfsBE2

Program terminated
{0} ok

```

Resetting ...

```

screen not found.
Can't open input device.
Keyboard not present. Using ttya for input and output.

```

```

Sun Enterprise 220R (2 X UltraSPARC-II 450MHz), No Keyboard
OpenBoot 3.23, 1024 MB memory installed, Serial #13116682.
Ethernet address 8:0:20:c8:25:a, Host ID: 80c8250a.

```

```

{0} ok boot -F failsafe -Z rpool/ROOT/zfsBE2
Boot device: /pci@1f,4000/scsi@3/disk@1,0 File and args: -F failsafe -Z

```

```
rpool/ROOT/zfsBE2
SunOS Release 5.10
Copyright 1983-2008 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
Configuring /dev
Searching for installed OS instances...
```

```
ROOT/zfsBE2 was found on rpool.
Do you wish to have it mounted read-write on /a? [y,n,?] y
mounting rpool on /a
```

Starting shell.

```
#
#
#
# zpool list
NAME      SIZE  USED  AVAIL    CAP  HEALTH  ALTROOT
rpool    16.8G  6.26G  10.5G   37%  ONLINE  /a
#
# zpool status
pool: rpool
state: ONLINE
scrub: none requested
config:
```

NAME	STATE	READ	WRITE	CKSUM
rpool	ONLINE	0	0	0
c0t1d0s0	ONLINE	0	0	0

errors: No known data errors

```
#
# df -h
Filesystem      size  used  avail  capacity  Mounted on
/ramdisk-root:a 163M  153M   0K    100%     /
/devices        0K    0K    0K     0%     /devices
/dev            0K    0K    0K     0%     /dev
ctfs            0K    0K    0K     0%     /system/contract
proc            0K    0K    0K     0%     /proc
mnttab          0K    0K    0K     0%     /etc/mnttab
swap            601M  344K  601M   1%     /etc/svc/volatile
objfs           0K    0K    0K     0%     /system/object
sharefs         0K    0K    0K     0%     /etc/dfs/sharetab
swap            602M  1.4M  601M   1%     /tmp
/tmp/root/etc   602M  1.4M  601M   1%     /.tmp_proto/root/etc
fd              0K    0K    0K     0%     /dev/fd
rpool/ROOT/zfsBE2 16G  5.7G  9.8G   37%     /a
rpool/export    16G  20K   9.8G   1%     /a/export
rpool/export/home 16G  18K   9.8G   1%     /a/export/home
rpool           16G  63K   9.8G   1%     /a/rpool
```

Booting a SPARC Based System From the Network

You might need to boot a system from the network under the following conditions:

- When the system is first installed
- If the system won't boot from the local disk
- If the system is a diskless client

Two network configuration boot strategies are available:

- Reverse Address Resolution Protocol (RARP) and ONC+ RPC Bootparams Protocol
- Dynamic Host Configuration Protocol (DHCP)

For network devices, the process for booting over a local area network (LAN) and booting over a wide area network (WAN) is slightly different. In both network boot scenarios, the PROM downloads the booter from a boot server or an installation server, which is `inetboot` in this case.

When booting over a (LAN), the firmware uses RARP and BOOTP or DHCP to discover the boot or installation server. TFTP is then used to download the booter, which is `inetboot` in this case.

When booting over a WAN, the firmware uses either DHCP or NVRAM properties to discover the installation server, the router, and the proxies that are required for the system to boot from the network. The protocol that is used to download the booter is HTTP. In addition, the booter's signature might be checked with a predefined private key.

▼ SPARC: How to Boot a System From the Network

Any system can boot from the network, if a boot server is available. You might want to boot a stand-alone system from the network if the system cannot boot from the local disk. For information on changing or resetting the default boot device, see [“SPARC: How to Change the Default Boot Device by Using the Boot PROM” on page 207](#).

Two network configuration boot strategies are available on sun-4u systems:

- RARP – Reverse Address Resolution Protocol and ONC+ RPC Bootparams Protocol
- DHCP – Dynamic Host Configuration Protocol

The default network boot strategy is set to RARP. You can use either protocol, depending on whether a RARP boot server or a DHCP boot server is available on your network.

Note – Sun Ultra systems must have at least PROM version 3.25.*nn* to use the DHCP network boot strategy. For information on determining your PROM version, see [“SPARC: How to Find the PROM Revision Number for a System”](#) on page 205.

If both protocols are available, you can temporarily specify which protocol to use in the boot command. Or, you can save the network boot strategy across system reboots at the PROM level by setting up an NVRAM alias. The following example uses the `nvalias` command to set up a network device alias for booting with DHCP by default on a Sun Ultra 10 system.

```
ok nvalias net /pci@1f,4000/network@1,1:dhcp
```

As a result, when you type `boot net`, the system boots by using DHCP.

Note – Do not use the `nvalias` command to modify the `NVRAMRC` file, unless you are very familiar with the syntax of this command and the `nvunalias` command. For information about using these commands, see the *OpenBoot 3.x Command Reference Manual*.

Before You Begin To boot with either protocol successfully, you must have already set up a RARP or DHCP boot server on your network.

- 1 **If necessary, shut down the system.**
- 2 **Determine the method for booting from the network, then select one of the following:**

a. Boot the system from the network by using the DHCP strategy.

```
ok boot net[:dhcp]
```

If you have changed the PROM setting to boot with DHCP by default, as in the preceding `nvalias` example, you only have to specify `boot net`.

b. Boot the system from the network by using the RARP strategy.

```
ok boot net[:rarp]
```

Because RARP is the default network boot strategy, you only have to specify `boot net : rarp` if you have changed the PROM value to boot DHCP.

Booting an x86 Based System by Using GRUB (Task Map)

Task	Description	For Instructions
Boot an x86 based system to run level 3, multiuser level.	Use this boot method to bring the system back to multiuser level after shutting down the system or performing a system hardware maintenance task.	“x86: How to Boot a System to Run Level 3 (Multiuser)” on page 246
Boot an x86 based system in single-user mode.	Use this boot method to perform a system maintenance task, such as backing up a file system.	“x86: How to Boot a System to Run Level S (Single-User Level)” on page 247
Boot an x86 based system interactively.	Use this boot method after making temporary changes to a system file or the kernel for testing purposes.	“x86: How to Boot a System Interactively” on page 249
Display a list of bootable ZFS datasets on an x86 based system.	Use one of the following methods to display the available BEs on an x86 based system that has a ZFS root file system: <ul style="list-style-type: none"> ■ <code>lustatus</code> ■ <code>bootadm list-menu</code> 	“x86: How to Display a List of the Available ZFS Boot Environments” on page 252
Boot an x86 based system from a ZFS root file system.	If you install or upgrade your system to an Oracle Solaris release that supports a ZFS boot loader, the GRUB menu entry for the default ZFS BE contains the <code>-B \$ZFS-BOOTFS</code> boot argument by default. The system boots automatically from ZFS. Note – This option is supported <i>only</i> for boot devices that contain a ZFS pool.	“x86: How to Boot From a Specified ZFS Root File System” on page 253
Boot an x86 based system in failsafe mode.	Use this procedure to boot the failsafe archive on an x86 based system. Then, run the <code>bootadm</code> command to update the boot archive.	“How to Boot an x86 Based System in Failsafe Mode” on page 256
Boot an x86 based system in failsafe mode to forcibly update a corrupt boot archive.	Use this procedure in cases where the boot archive is corrupt, and the system refuses to boot normally, or you are not prompted to update an inconsistent boot archive.	“x86: How to Boot in Failsafe Mode to Forcibly Update a Corrupt Boot Archive” on page 258
Boot an x86 based system from the network by using GRUB.	Use this method to boot a PXE or non-PXE device from the network with the default network configuration strategy. This method is also used for booting a diskless client.	“x86: How to Perform a GRUB Based Boot From the Network” on page 263

▼ x86: How to Boot a System to Run Level 3 (Multiuser)

Use this procedure to boot a system that is currently at run level 0 to run level 3.

1 Reboot the system.

```
# reboot
```

If the system displays the Press any key to reboot prompt, press any key to reboot the system.

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

When the boot sequence begins, the GRUB menu is displayed.

2 When the GRUB menu is displayed, press Enter to boot the default OS instance.

If you do not choose an entry within 10 seconds, the system automatically boots to run level 3.

The login prompt is displayed when the boot process has finished successfully.

3 Log in to the system.

```
hostname console login:
```

4 Verify that the system booted to run level 3.

```
# who -r
system% who -r
.          run-level 3  Mar  2 09:44    3      0  S
```

Example 12-9 x86: Booting a System To Run Level 3 (Multiuser Level)

```
# reboot
```

```
Jul 24 11:29:52 bearskin reboot: rebooted by root
syncing file systems... done
rebooting...
```

```
Adaptec AIC-7899 SCSI BIOS v2.5754
(c) 2000 Adaptec, Inc. All Rights Reserved.
```

```
Press <Ctrl><A> for SCSISelect(TM) Utility!
```

```
Ch B, SCSI ID: 0 SEAGATE ST336607LSUN36G 160
```

```
GNU GRUB version 0.95 (637K lower / 2096064K upper memory)
```

```
=====
Solaris 10 10/08 s10x_u6wos_03 X86
Solaris failsafe
```

```
=====
Use the and keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
```

commands before booting, or 'c' for a command-line.

```
SunOS Release 5.10 Version Generic_144500-10 64-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
Hostname: pups
NIS domain name is ...sfbay.sun.com
Reading ZFS config: done.
Mounting ZFS filesystems: (5/5)
```

pups console login:

```
# who -r
.          run-level 3  Jul 24 11:31    3      0  S
```

▼ x86: How to Boot a System to Run Level S (Single-User Level)

Use this procedure to boot a system that is at run level 0 to run level S. The single-user level is used for performing system maintenance.

Note – This procedure can be used for all GRUB implementations. However, the boot entries in the GRUB main menu vary, depending on which Oracle Solaris release you are running.

For a description of all of the kernel options that you can specify at boot time, see [“x86: Modifying Boot Behavior by Editing the GRUB Menu at Boot Time” on page 214](#).

1 Reboot the system.

```
# reboot
```

If the system displays the Press any key to reboot prompt, press any key to reboot the system.

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

When the boot sequence begins, the GRUB menu is displayed.

2 When the GRUB main menu is displayed, type e to edit the GRUB menu.

3 Depending on the release you are running, use the arrow keys to choose the kernel or kernel\$ line.

If you cannot use the arrow keys, use the caret key (^) key to scroll up and the letter v key to scroll down.

4 Type e again to edit the boot entry.

From here, you can add options and arguments to the `kernel` or `kernel$` line.

5 To boot the system in single-user mode, type -s at the end of the boot entry line, then press Return to go back to the previous screen.

▪ **To specify other boot behaviors, replace the -s option with the appropriate boot option.**

The following alternate boot behaviors can be specified in this manner:

- Perform a reconfiguration boot
- Boot a 64-bit capable system in 32-bit mode
- Boot the system with the kernel debugger
- Redirect the console

For more information, see the [boot\(1M\)](#) man page.

6 To boot the system in single-user mode, type b.

7 When prompted, type the root password.

8 Verify that the system is at run level S.

```
# who -r
.          run-level S  Jun 13 11:07      S      0  0
```

9 Perform the system maintenance task that required the run level change to S.

10 After you complete the system maintenance task, reboot the system.

Example 12–10 x86: Booting a System in Single-User Mode

```
# reboot
Jul  2 14:30:01 pups reboot: initiated by root on /dev/console
syncing files...

Press <Ctrl><A> forPSCSISelect(TM) Utility!

GNU GRUB  version 0.95  (637K lower / 2096064K upper memory)

=====
Solaris 10 10/08 s10x_u6wos_03 X86
Solaris failsafe

=====
Use the  and  keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.
=====

GNU GRUB  version 0.95  (637K lower / 2096064K upper memory)
```



```

=====
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive
=====
      Use the  and  keys to select which entry is highlighted.
      Press 'b' to boot, 'e' to edit the selected command in the
      boot sequence, 'c' for a command-line, 'o' to open a new line
      after ('O' for before) the selected line, 'd' to remove the
      selected line, or escape to go back to the main menu.

[ Minimal BASH-like line editing is supported.  For the first word, TAB
lists possible command completions.  Anywhere else TAB lists the possible
completions of a device/filename.  ESC at any time exits. ]

grub edit> kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS -s

GNU GRUB  version 0.95  (637K lower / 2096064K upper memory)

=====
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS -s
module /platform/i86pc/boot_archive
=====
      Use the  and  keys to select which entry is highlighted.
      Press 'b' to boot, 'e' to edit the selected command in the
      boot sequence, 'c' for a command-line, 'o' to open a new line
      after ('O' for before) the selected line, 'd' to remove the
      selected line, or escape to go back to the main menu.
.
.
.
SunOS Release 5.10 Version Generic_144500-10 64-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
Booting to milestone "milestone/single-user:default".
Hostname: pups Requesting System Maintenance Mode SINGLE USER MODE
Root password for system maintenance (control-d to bypass):
single-user privilege assigned to /dev/console.
Entering System Maintenance Mode
Jul  2 14:41:48 su: 'su root' succeeded for root on /dev/console Sun Microsystems Inc.
# who -r
who -r      .          run-level S Jul  2 14:39      S      0  0 #

```

▼ x86: How to Boot a System Interactively

Use this procedure to boot a system, if you need to specify an alternate kernel or an alternate `/etc/system` file.

Before You Begin To specify an alternate `/etc/system` file when booting an x86 based system interactively by using the `boot -a` command, follow these steps:

- 1. Make backup copies of the `/etc/system` and the `boot/solaris/filelist.ramdisk` files.

- ```
cp /etc/system /etc/system.bak
cp /boot/solaris/filelist.ramdisk /boot/solaris/filelist.ramdisk.orig
```
- 2. Add the /etc/system.bak file name to the /boot/solaris/filelist.ramdisk file.

```
echo "etc/system.bak" >> /boot/solaris/filelist.ramdisk
```
  - 3. Update the boot archive.

```
bootadm update-archive -v
```

### 1 Reboot the system.

```
reboot
```

If the system displays the Press any key to reboot prompt, press any key to reboot the system.

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

When the boot sequence begins, the GRUB main menu is displayed.

### 2 To access the GRUB edit menu, type e.

### 3 Use the arrow keys to select the kernel or kernel\$ line.

### 4 Type e to edit the boot entry line.

### 5 Type -a to boot the system interactively, then, press Enter to return to the previous menu.

### 6 To boot the system interactively, type b.

### 7 Type a default directory for modules, or press Enter to accept the default.

```
Enter default directory for modules [/platform/i86pc/kernel /kernel /usr/kernel]:
```

### 8 Type an alternate system file name, *alternate-file*.

```
Name of system file [etc/system]: /etc/system.bak
```

Pressing Enter without providing an alternate file accepts the default.

Repair the damaged /etc/system file.

### 9 Reboot the system to run level 3.

## Example 12-11 x86: Booting a System Interactively

```
reboot
syncing file systems... done
rebooting...
```

```

GNU GRUB version 0.95 (637K lower / 2096064K upper memory)
=====
Solaris 10 10/08 s10x_u6wos_03 X86
Solaris failsafe
=====
 Use the and keys to select which entry is highlighted.
 Press enter to boot the selected OS, 'e' to edit the
 commands before booting, or 'c' for a command-line.
=====

```

```

GNU GRUB version 0.95 (637K lower / 2096064K upper memory)
=====
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive
=====

```

```

 Use the and keys to select which entry is highlighted.
 Press 'b' to boot, 'e' to edit the selected command in the
 boot sequence, 'c' for a command-line, 'o' to open a new line
 after ('O' for before) the selected line, 'd' to remove the
 selected line, or escape to go back to the main menu.

```

```

[Minimal BASH-like line editing is supported. For the first word, TAB
lists possible command completions. Anywhere else TAB lists the possible
completions of a device/filename. ESC at any time exits.]

```

```

grub edit> kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS -a
GNU GRUB version 0.95 (637K lower / 2096064K upper memory)

```

```

=====
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS -a
module /platform/i86pc/boot_archive
=====

```

```

.
.
.

```

```

Enter default directory for modules [/platform/i86pc/kernel /kernel /usr/kernel]:
Name of system file [/etc/system]: /etc/system.bak
SunOS Release 5.10 Version Generic_144500-10 64-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
Hostname: pups
NIS domain name is ...sfbay.sun.com
Reading ZFS config: done.
Mounting ZFS filesystems: (5/5)
pups console login:

```

## x86: Booting From a Specified ZFS Root File System on an x86 Based System

To support booting an Oracle Solaris ZFS root file system on the x86 platform, a new GRUB keyword, `$ZFS-BOOTFS`, has been introduced. If a root device contains a ZFS pool, this keyword is assigned a value, which is then passed to the kernel with the `-B` option. This option identifies which dataset to boot. If you install or upgrade your system with an Oracle Solaris release that supports a ZFS boot loader, the GRUB menu `.lst` file, as well as the GRUB boot menu, contain this information by default.

### ▼ x86: How to Display a List of the Available ZFS Boot Environments

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

#### 2 To display a list of available BEs on the system, type the following command:

```
~# bootadm list-menu
```

```
lustatus
```

Note that the `lustatus` command can also be used on SPARC based systems.

---

**Note** – If the following error is displayed when you run the `lustatus` command, it is an indication that a new installation was performed and that Solaris Live Upgrade was not used. Before any BEs can be acknowledged in the `lustatus` output, a new BE must be first created on the system.

```
lustatus
ERROR: No boot environments are configured on this system
ERROR: cannot determine list of all boot environment names
```

---

For more information about using Solaris Live Upgrade to migrate a UFS root file system to a ZFS root file system, see “[Migrating to a ZFS Root File System or Updating a ZFS Root File System \(Live Upgrade\)](#)” in *Oracle Solaris ZFS Administration Guide*.

#### Example 12–12 Displaying a List of Available ZFS Bootable Datasets by Using the `lustatus` Command

In this example, the output of the `lustatus` command shows the status of three ZFS bootable datasets. The default boot environment is `be1` and therefore cannot be deleted.

```
lustatus
Boot Environment Is Active Active Can Copy
Name Complete Now On Reboot Delete Status

s10s_nbu6wos yes no no yes -
zfs2BE yes yes yes no -
zfsbe3 no no no yes -
#
```

If the BE has been created and is bootable, a “yes” appears in the Is Complete column. If a BE has been created, but is not yet activated, a “no” appears in this column. To activate a BE, use the `luactivate` command. Run the `lustatus` command afterwards to verify that the BE was successfully activated.

For more information see the `lustatus(1M)` and the `luactivate(1M)` man pages.

## ▼ x86: How to Boot From a Specified ZFS Root File System

This procedure describes how to boot from a ZFS root file system on an x86 system that supports a ZFS boot loader.

Note that if you install or upgrade your system to an Oracle Solaris release that supports a ZFS boot loader, the GRUB menu entry contains the `-B $ZFS-BOOTFS` boot argument by default, so the system boots from ZFS without requiring any additional boot arguments.

### 1 Reboot the system.

```
reboot
```

If the system displays the Press any key to reboot prompt, press any key to reboot the system.

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

When the boot sequence begins, the GRUB main menu is displayed. If the default boot entry is a ZFS file system menu is similar to the following:

```
GNU GRUB version 0.95 (637K lower / 3144640K upper memory)
+-----+
| be1
| be1 failsafe
| be3
| be3 failsafe
| be2
| be2 failfafa
+-----+
Use the ^ and v keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.
```

**2 When the GRUB menu is displayed, press Enter to boot the default OS instance.**

If you do not choose an entry within 10 seconds, the system automatically boots to run level 3.

**3 To boot another BE, use the arrow keys to highlight the specified boot entry.****4 Type b to boot this entry or e to edit the entry.**

For more information about GRUB menu entries at boot time, see “[x86: How to Modify Boot Behavior by Editing the GRUB Menu at Boot Time](#)” on page 217.

**Example 12–13 x86: Activating a New Boot Environment on an x86 Based System**

This example shows the steps that are followed to activate a boot environment, `be10`, on a system. Note that the `lustatus` command is run first, to determine which BEs on the system are active and which BEs require activation.

```
lustatus
Boot Environment Is Active Active Can Copy
Name Complete Now On Reboot Delete Status

be1 yes yes yes no
be10 yes no no yes

luactivate be10
System has findroot enabled GRUB Generating boot-sign, partition and slice
information for PBE <be1>
WARNING: The following file s have change on both the current boot environment
<be1> zone <global> and the boot environment to be activitate <be10>
/etc/zfs/zpool.cache
INFORMATION: The files listed above are in conflict between the current
boot environment <be1> zone <global> and the boot environment to be
activated <be10>. These files will not be automatically synchronized from
the current boot environment <be1> when boot environment <be10> is activated.

Setting failsafe console to <ttyb>
Generating boot-sign for ABE <be10>
Generating partition and slice information for ABE <be10>
Copied boot menu from top level dataset.
Generating direct boot menu entries for PBE.
Generating direct boot menu entries for ABE.
Disabling splashimage
Current GRUB menu default setting is not valid
title Solaris bootenv rc
No more bootadm entries. Deletion of bootadm entries is complete.
GRUB menu default setting is unchanged
Done eliding bootadm entries.

The target boot environment has been activated. It will be used when you
reboot. NOTE: You MUST NOT USE the reboot, halt, or uadmin commands. You
MUST USE either the init or the shutdown command when you reboot. If you
do not use either init or shutdown, the system will not boot using the
target BE.
```

```

', '

reboot
May 30 09:52:32 pups reboot: initiated by root on /dev/console
syncing file systems... done
rebooting...

CE SDRAM BIOS P/N GR-xlint.007-4.330
*

BIOS Lan-Console 2.0
Copyright (C) 1999-2001 Intel Corporation
.
.
GNU GRUB version 0.95 (637K lower / 3144640K upper memory)
+-----+
| be1
| be1 failsafe
| be10
| be10 failsafe
+-----+
 Use the ^ and v keys to select which entry is highlighted.
 Press enter to boot the selected OS, 'e' to edit the
 commands before booting, or 'c' for a command-line.

SunOS Release 5.10 Version Generic_144500-10 64-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.

Hostname: pups
NIS domain name is sunsoft.eng.sun.com
Reading ZFS config: done.
Mounting ZFS filesystems: (8/8)

pups console login:
lustatus
Boot Environment Is Active Active Can Copy
Name Complete Now On Reboot Delete Status

be1 yes yes yes no
be10 yes yes yes no
#

```

## Booting an x86 Based System in Failsafe Mode

To boot an x86 based system in failsafe mode, select the failsafe boot entry when the GRUB menu is displayed during a system boot. During the failsafe boot procedure, when prompted by the system, type `y` to update the primary boot archive.

Failsafe booting is also supported on systems that are booted from ZFS. When booting from a UFS-rooted BE, each BE has its own failsafe archive. The failsafe archive is located where the root file system is located, as is the case with a ZFS-rooted BE. On x86 based systems, each

failsafe archive has an entry in the pool-wide GRUB menu. The default failsafe archive is the archive that is in the default bootable file system. The default bootable file system (dataset) is indicated by the value of the pool's `boot fs` property.

For more information about boot archive recovery, see the [Chapter 13, “Managing the Oracle Solaris Boot Archives \(Tasks\)”](#).

## ▼ How to Boot an x86 Based System in Failsafe Mode

**Note** – The GRUB failsafe interaction in some Oracle Solaris releases prompts you to update the boot archive, regardless of whether any inconsistent boot archive are detected. In this release, the system only prompts you to update the boot archive if an inconsistent boot archive is detected.

- 1 **Stop the system by using one of the methods described in the procedure, “x86: How to Stop a System for Recovery Purposes” on page 284.**
- 2 **If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the Reset button at this prompt. Or, you can use the power switch to reboot the system.

When the boot sequence begins, the GRUB menu is displayed.

```
GNU GRUB version 0.95 (637K lower / 3144640K upper memory)
+-----+
| be1
| be1 failsafe
| be3
| be3 failsafe
| be2
| be2 failfafa
+-----+
 Use the ^ and v keys to select which entry is highlighted.
 Press enter to boot the selected OS, 'e' to edit the
 commands before booting, or 'c' for a command-line.
```

**Note** – The GRUB menu that is displayed varies, depending on which Oracle Solaris release you are running.

- 3 **Use the arrow keys to navigate the GRUB menu to select a failsafe entry.**



**4 Press Return to boot the failsafe archive.**

The system searches for installed OS instances. If an inconsistent boot archive is detected, a message similar to the following is displayed:

```
Searching for installed OS instances...
```

```
An out of sync boot archive was detected on /dev/dsk/c0t0d0s0.
The boot archive is a cache of files used during boot and
should be kept in sync to ensure proper system operation.
```

```
Do you wish to automatically update this boot archive? [y,n,?]
```

**5 Type y to update the boot archive.**

If multiple inconsistent boot archives are detected, the system will prompt you to type y to update each inconsistent boot archive.

For each archive that is updated successfully, the following message is displayed:

```
Updating boot archive on /dev/dsk/c0t0d0s0.
```

```
The boot archive on /dev/dsk/c0t0d0s0 was updated successfully.
```

After the boot archive is updated, the system searches again for all installed OS instances, then prompts you to select a device to mount on /a. Note that this same message is displayed when the system first boots if no inconsistent boot archives are detected.

```
Searching for installed OS instances...
```

```
Multiple OS instances were found. To check and mount one of them
read-write under /a, select it from the following list. To not mount
any, select 'q'.
```

```
1 pool10:13292304648356142148 ROOT/be10
2 rpool:14465159259155950256 ROOT/be01
```

```
Please select a device to be mounted (q for none) [?,??,q]:
```

- **If you choose not to mount a device, type q to continue to boot process.**
- **If you choose to mount a device, follow these steps:**
  - a. **Type the number of the device, then press Return.**  
The system mounts the device on /a, and returns you to a shell prompt.
  - b. **Repair the critical system resource.**
  - c. **When you are done repairing the critical system resource, unmount the device.**  
# `umount /a`
  - d. **Reboot the system.**  
# `reboot`

## ▼ x86: How to Boot in Failsafe Mode to Forcibly Update a Corrupt Boot Archive

This procedure shows how to rebuild an inconsistent or corrupt boot archive in the event you are not prompted by the system to update the boot archive the system, or in the event of a system hang or looping sequence occurs.

- 1 **Stop the system by using one of the methods that are described in the procedure, “x86: How to Stop a System for Recovery Purposes” on page 284.**

- 2 **Reboot the system.**

# **reboot**

If the system displays the Press any key to reboot prompt, press any key to reboot the system.

You can also use the Reset button at this prompt.

When the boot sequence begins, the GRUB menu is displayed.

```
+-----+
| Solaris 10.1... X86 |
| Solaris failsafe |
| |
+-----+
Use the and keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.
```

---

**Note** – The contents of the GRUB menus vary, depending on the Solaris release you are running.

---

- 3 **Use the arrow keys to navigate the GRUB menu to select the failsafe entry.**

- 4 **Press Return to boot the failsafe archive.**

If any boot archives are out of date, a message t similar to the following is displayed:

Searching for installed OS instances...

```
An out of sync boot archive was detected on /dev/dsk/c0t0d0s0.
The boot archive is a cache of files used during boot and
should be kept in sync to ensure proper system operation.
```

```
Do you wish to automatically update this boot archive? [y,n,?]
```

**5 Type y, then press Enter to update the inconsistent boot archive.**

The system displays the following message:

```
Updating boot archive on /dev/dsk/c0t0d0s0.
```

```
The boot archive on /dev/dsk/c0t0d0s0 was updated successfully.
```

If no inconsistent boot archives are found, a message similar to the following is displayed:

```
Searching for installed OS instances...
```

```
Solaris 10.1... X86 was found on /dev/dsk/c0t0d0s0.
```

```
Do you wish to have it mounted read-write on /a? [y,n,?]
```

This message is also displayed after any inconsistent boot archives are updated successfully.

**6 Mount the device that contains the corrupt boot archive on /a by typing the corresponding number of the device, then pressing Enter..**


---

**Note** – If any inconsistent boot archives were updated in the previous step, the device is already mounted on /a.

---

**7 To forcibly update the corrupt boot archive, type:**

```
bootadm update-archive -f -R /a
```

**8 Unmount the device.**

```
umount /a
```

**9 Reboot the system.**

```
reboot
```

**Example 12–14 x86: Booting in Failsafe Mode to Forcibly Update a Corrupt Boot Archive**

This example shows how to boot the failsafe archive to forcibly update a corrupt boot archive.

```
GNU GRUB version 0.95 (635K lower / 523200K upper memory)
```

```
+-----+
| Solaris 10 1/06 s10x_u1wos_19a X86 |
| >Solaris failsafe< |
| |
+-----+
```

```
Use the and keys to select which entry is highlighted.
Press enter to boot the selected OS, 'e' to edit the
commands before booting, or 'c' for a command-line.
```

```
SunOS Release 5.10
```

```
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
```

```
Use is subject to license terms.
```

```
Booting to milestone "milestone/single-user:default".
```

```

Configuring devices.
Searching for installed OS instances...

Multiple OS instances were found. To check and mount one of them
read-write under /a, select it from the following list. To not mount
any, select 'q'.

 1 /dev/dsk/c0t0d0s0 Solaris 10 1/06 s10x_ulwos_19a X86
 2 /dev/dsk/c0t1d0s0 Solaris 10 5/08 X86

Please select a device to be mounted (q for none) [?,??,q]: 1
mounting /dev/dsk/c0t0d0s0 on /a

Starting shell.
rm /a/platform/i86pc/boot_archive
bootadm update-archive -f -R /a
Creating boot_archive for /a
updating /a/platform/i86pc/amd64/boot_archive
updating /a/platform/i86pc/boot_archive
umount /a
reboot
syncing file systems... done
rebooting...
.
.
.

```

## Booting an x86 Based System From the Network

This section describes the requirements and warnings for performing a GRUB based boot from the network.

Any system can boot from the network, if a boot server is available. You might need to boot a stand-alone system from the network for recovery purposes if the system cannot boot from the local disk. You can boot an x86 based system directly from a network that supports the PXE network boot protocol.

---

**Note** – The PXE network boot is available only for devices that implement the Intel Preboot Execution Environment specification.

---

The default network boot strategy that is used for a GRUB based PXE network boot is DHCP. For non-PXE devices, you can use either the DHCP or the RARP boot strategy. The strategy that you use depends on which type of boot server is available on your network. If no PXE or DHCP server is available, you can load GRUB from a diskette, a CD-ROM, or a local disk.

To perform a GRUB based network boot, a DHCP server that is configured for PXE clients is required. A boot server that provides tftp service is also required. The DHCP server supplies the information that the client needs to configure its network interface.

The DHCP server must be able to respond to the DHCP classes, `PXEClient` and `GRUBClient` with the following information:

- IP address of the file server
- Name of the boot file (`pxegrub`)

The sequence for performing a PXE network boot of the Oracle Solaris OS is as follows:

1. The BIOS is configured to boot from a network interface.
2. The BIOS sends a DHCP request.
3. The DHCP server replies with the server address and the name of the boot file.
4. The BIOS downloads `pxegrub` by using `tftp` and executes `pxegrub`.
5. The system downloads a GRUB menu file by using `tftp`.  
This file displays the boot menu entries that are available.
6. After you select a menu entry, the system begins to load the Oracle Solaris OS.

See “[How to Set Up a Network Configuration Server](#)” in *System Administration Guide: IP Services* for more information.

Running the `add_install_client` command creates the `/tftpboot_01ethernet-address` file. This file is linked to `pxegrub` and the `/tftpboot/menu.lst.01ethernet-address` file. The `/tftpboot/menu.lst.01ethernet-address` file is the GRUB menu file. If this file does not exist, then `pxegrub` reverts to using DHCP Option 150, if this option is specified, or the `/tftpboot/boot/grub/menu.lst` file. Typically, a single system is set up to serve both functions. In this instance, the `add_install_client` command sets up the `/tftpboot` file with the correct `pxegrub` menu file and the Oracle Solaris files. DHCP service is handled separately by using the `add_install_client` command. The setup only needs to be completed once per client. See “[x86: About DHCP Macros](#)” on page 261 and “[x86: How to Perform a GRUB Based Boot From the Network](#)” on page 263 for more information.

## x86: About DHCP Macros

When you add clients with the `add_install_client -d` script on the install server, the script reports DHCP configuration information to standard output. You can use this information when you create the options and macros that are needed to pass network installation information to clients.

To install DHCP clients with a DHCP server over the network, you must create DHCP options. This information is needed to install the Oracle Solaris OS.

When a client sends a DHCP request, the server must have the following client information:

- Client's ID, which is typically the Ethernet address
- Class of the client request

- Subnet on which the client resides

The DHCP server forms a response. This response is based on the following *macros*, which matches the client request:

|                      |                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>class macro</b>   | The class macro is based on a <i>class string</i> that is contained in the DHCP request. On x86 based systems, the BIOS already makes a DHCP request with the class <code>PXEClient:Arch:00000:UNDI:002001</code> . If a macro by this name is defined in the DHCP server configuration, then the macro content is sent to the x86 based clients.                                                   |
| <b>network macro</b> | The network macro is named by the IP address of the subnet that the client resides on. If the macro <code>129.146.87.0</code> is defined on the DHCP server, the macro content is sent to all clients on that subnet. The macro content is sent, regardless of the class of the request. If an option is defined in both the class macro and the network macro, the network macro takes precedence. |
| <b>IP macro</b>      | The IP macro is named by an IP address. This macro is rarely used.                                                                                                                                                                                                                                                                                                                                  |
| <b>client macro</b>  | The client macro is named by the client type (01 for Ethernet) and the MAC address of the client, in uppercase letters. For a client with the Ethernet address <code>0:0:39:fc:f2:ef</code> , the corresponding macro name is <code>01000039FCEF</code> . Note the absence of colons in the client macro.                                                                                           |

For example, for a client on the subnet `192.168.100.0`, with the Ethernet address `0:0:39:fc:f2:ef`, making a DHCP request of class `PXEClient`, the DHCP server has the following matching macro:

```
PXEClient
 BootSrvA: 192.168.100.0
 BootFile: pxegrub
129.146.87.0
 Router: 129.146.87.1
 NISdmain: sunsoft.eng.sun.com
01000039FCEF
 BootFile: 01000039FCEF
The actual DHCP response will be
 BootSrvA: 192.168.100.0
 BootFile: 01000039FCEF
 Router: 129.146.87.1
 NISdmain: sunsoft.eng.sun.com
```

Note that the `BootFile` in the client macro overrides the `BootFile` in the class macro.

For more detailed information, see “[Preconfiguring System Configuration Information With the DHCP Service \(Tasks\)](#)” in *Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations*.

## ▼ x86: How to Perform a GRUB Based Boot From the Network

To perform a GRUB based network boot, a DHCP server that is configured for PXE clients is required. A boot server that provides tftp service is also required. The DHCP server must be able respond to the DHCP classes, PXEClient and GRUBClient, to obtain the IP address of the file server and the boot file (pxegrub). By default, the menu file is /tftpboot/menu.lst.01ethernet-address. If this file does not exist, then pxegrub reverts to DHCP Option 150, if this option is specified, or the /tftpboot/boot/grub/menu.lst file.

If you are booting the system from the Solaris software media, the system boots automatically.

**Before You Begin** Before performing a network boot on an x86 based system with GRUB, do the following:

- Run the appropriate commands on the installation server to enable the system to boot from the network.
- Add the client system as an install client.

See [Chapter 4, “Installing From the Network \(Overview\)”](#) in *Oracle Solaris 10 8/11 Installation Guide: Network-Based Installations* for more information.

### 1 On the DHCP server, create a client macro for the DHCP service with the following two options:

- BootSrvA: *svr-addr*
- BootFile: *client-macro*

Note that you must have superuser privileges on the DHCP server to run the dhtadm command.

where *svr-addr* is the IP address of the server, and *client-macro* is named by the client's Ethernet type (01) and the MAC address, in uppercase letters. This number is also the name of the file that is used in the /tftpboot directory on the installation server.

---

**Note** – The notation for the *client-macro* should not contain any colons.

---

You can create the client macro from the DHCP GUI or from command-line interface.

To create the client macro from the command-line, type:

```
dhtadm -[MA] -m client macro -d
":BootFile=client-macro:BootSrvA=svr-addr:"
```

### 2 Reboot the system.

- 3 Instruct the BIOS to boot from the network.
  - If your system uses a specific keystroke sequence to boot from the network, type the keystrokes when the BIOS screen is displayed.
  - If you need to manually modify the BIOS settings to boot from the network, type the keystroke sequence to access the BIOS setup utility. Then, modify the boot priority to boot from the network.
- 4 When the GRUB menu is displayed, select the network installation image that you want to install.

## Accelerating the Reboot Process on the SPARC Platform (Task Map)

| Task                                                          | Description                                                                                                                                                                                                                                                                                                                                               | For Instructions                                                                    |
|---------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Initiate a fast reboot of a SPARC based system.               | <p>If the Fast Reboot feature is not enabled on a SPARC based system, use the <code>reboot</code> command with the <code>-f</code> option to initiate a fast reboot of the system.</p> <p>If the Fast Reboot feature has been enabled, use either the <code>reboot</code> or the <code>init 6</code> command to initiate a fast reboot of the system.</p> | <a href="#">“How to Initiate a Fast Reboot of a SPARC Based System” on page 265</a> |
| Perform a standard reboot of a SPARC based system.            | Use the <code>reboot</code> command with the <code>-p</code> option to perform a standard reboot of the system.                                                                                                                                                                                                                                           | <a href="#">“Performing a Standard Reboot of a SPARC Based System” on page 265</a>  |
| Enable fast reboot behavior by default on the SPARC platform. | On the SPARC platform, Fast Reboot behavior is disabled by default. You can configure the <code>boot-config</code> service to perform a fast reboot of a SPARC based system by default.                                                                                                                                                                   | <a href="#">“Managing the Boot Configuration Service” on page 265</a>               |



# Initiating a Fast Reboot of a SPARC Based System

The Fast Reboot feature of Oracle Solaris is supported on the SPARC platform. This section describes typical tasks that you might be required to perform.

## ▼ How to Initiate a Fast Reboot of a SPARC Based System

Use the following procedure to initiate fast reboot of a SPARC based system when the `config/fastreboot_default` property of the `boot-config` service is set to `false`, which is the default behavior. To change the default behavior of the Fast Reboot feature so that a fast reboot is automatically performed when the system reboots, see [“Managing the Boot Configuration Service” on page 265](#).

- 1 Assume the root role.
- 2 Initiate a fast reboot of the system by typing the following command:

```
reboot -f
```

## Performing a Standard Reboot of a SPARC Based System

In some situations, such as when you are booting a system from the network, certain POST tests must be performed during the boot process. To reboot a SPARC based system without skipping any POST tests without having to disable the Fast Reboot default behavior, use the `-p` option with the `reboot` command, as shown in the following example:

```
reboot -p
```

## Managing the Boot Configuration Service

The `fastreboot_default` property of the `boot-config` service enables an automatic fast reboot of the system when either the `reboot` or the `init 6` command is used. By default, this property's value is set to `false` on a SPARC based system.

The default behavior for this property can be configured by using the `svccfg` and `svcadm` commands. The following example shows how to set the property's value to `true` so that a fast reboot is initiated by default on the SPARC platform.

```
svccfg -s "system/boot-config:default" setprop config/fastreboot_default=true
svcadm refresh svc:/system/boot-config:default
```

For information about managing the boot configuration service through the SMF, see the [`svcadm\(1M\)`](#) and the [`svccfg\(1M\)`](#) man pages.



# Managing the Oracle Solaris Boot Archives (Tasks)

---

This chapter describes how the Oracle Solaris boot archives are managed. Procedures for using the `bootadm` command are described in detail.

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

- “Managing the Oracle Solaris Boot Archives (Task Map)” on page 267
- “Description of the Oracle Solaris Boot Archives” on page 268
- “Managing the boot - archive Service” on page 269
- “Automatic Boot Archive Recovery” on page 270
- “Using the `bootadm` Command to Manage the Boot Archives” on page 272

For overview information about the boot process, see [Chapter 9, “Shutting Down and Booting a System \(Overview\)”](#). For step-by-step instructions on booting a system, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\)”](#).

## Managing the Oracle Solaris Boot Archives (Task Map)

TABLE 13-1 Managing the Boot Archives (Task Map)

| Task                               | Description                                                                                                                                                                                                                                       | For Instructions                                                  |
|------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------|
| Manage the boot - archive service. | The boot - archive service is controlled by the Service Management Facility (SMF). Use the <code>svcadm</code> command to enable and disable services. Use the <code>svcs</code> command to verify whether the boot - archive service is running. | <a href="#">“Managing the boot - archive Service” on page 269</a> |

TABLE 13-1 Managing the Boot Archives (Task Map) (Continued)

| Task                                                                                                       | Description                                                                                                                                                                                   | For Instructions                                                                                                                           |
|------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
| <b>x86:</b> Clear a boot archive update failure by using the <code>auto-reboot-safe</code> property.       | Use this procedure in cases where the boot archive update on an x86 based system fails because the <code>auto-reboot-safe</code> property is set to <code>false</code> .                      | <a href="#">“x86: How to Clear Automatic Boot Archive Update Failures by Using the <code>auto-reboot-safe</code> Property” on page 270</a> |
| Clear a boot archive update failure by using the <code>bootadm</code> command.                             | Use this procedure to manually clear boot archive update failures on the SPARC platform, and on the x86 platform, if the <code>auto-reboot-safe</code> property is set to <code>true</code> . | <a href="#">“How to Clear Automatic Boot Archive Update Failures by Using the <code>bootadm</code> Command” on page 271</a>                |
| Manually update the boot archives by using the <code>bootadm</code> command.                               | Use the <code>bootadm update-archive</code> command to manually update the boot archive.                                                                                                      | <a href="#">“How to Manually Update the Boot Archive by Using the <code>bootadm</code> Command” on page 272</a>                            |
| Manually update the boot archive on a system with a Solaris Volume Manager (SVM) metadevice mirrored root. | On systems that use a metadevice mirror, you must manually mount the device before updating the boot archive by using the <code>bootadm</code> command.                                       | <a href="#">“How to Manually Update the Boot Archive on a Solaris Volume Manager RAID-1 (Mirrored) Root Partition” on page 273</a>         |
| List the contents of the boot archives by using the <code>bootadm</code> command.                          | Use the <code>bootadm list-archive</code> command to list the contents of the boot archive.                                                                                                   | <a href="#">“How to List Contents of the Boot Archive” on page 275</a>                                                                     |
| <b>x86:</b> Locate the active GRUB menu by using the <code>bootadm</code> command.                         | Use the <code>bootadm list-menu</code> command to determine the location of the active GRUB menu.                                                                                             | <a href="#">“x86: How to Locate the Active GRUB Menu and List Current Menu Entries” on page 275</a>                                        |
| <b>x86:</b> Set the default boot entry in the GRUB menu by using the <code>bootadm</code> command.         | Use the <code>bootadm set-menu</code> command to set the default boot entry in the GRUB menu.                                                                                                 | <a href="#">“x86: How to Set the Default Boot Entry for the Active GRUB Menu” on page 276</a>                                              |

## Description of the Oracle Solaris Boot Archives

When you install Oracle Solaris OS on a system, the `bootadm` command creates one primary boot archive and one failsafe archive.

A *primary boot archive* is a subset of a root (`/`) file system. This boot archive contains all of the kernel modules, `driver.conf` files, and some configuration files. These files are located in the `/etc` directory. The files in the boot archive are read by the kernel before the root (`/`) file system is mounted. After the root (`/`) file system is mounted, the boot archive is discarded by the kernel from memory. Then, file I/O is performed against the root device.

The files that make up the SPARC boot archives are located in the `/platform` directory.

The contents of this directory are divided into three groups of files:

- Files that are required for a sun4u boot archive
- Files that are required for a sun4v boot archive
- Files that are required for a sun4us boot archive

The files that make up the x86 boot archives are located in the `/platform/i86pc` directory.

To list the files and directories that are included in the boot archives, use the `bootadm list-archive` command.

If any files in the archive are updated, the boot archive must be rebuilt. For modifications to take effect, the rebuild of the archive must take place before the next system reboot.

The *failsafe* boot archive is the second type of archive that is created when you install the Solaris OS.

A failsafe boot archive has the following benefits and characteristics:

- Is self-sufficient
- Can boot on its own
- Is created by default during installation of the OS
- Requires no maintenance

For more information about booting a system in failsafe mode, see “[Booting a SPARC Based System in Failsafe Mode](#)” on page 238 and “[Booting an x86 Based System in Failsafe Mode](#)” on page 255.

## Managing the boot-archive Service

The boot-archive service is controlled by the Service Management Facility (SMF). The boot-archive service instance is `svc:/system/boot-archive:default`. The `svcadm` command is used to enable and disable services.

To determine if the boot-archive service is running, use the `svcs` command.

For more information, see the [svcadm\(1M\)](#) and the [svcs\(1\)](#) man pages.

### ▼ How to Enable or Disable the boot-archive Service

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

#### 2 To enable or disable the boot-archive service, type:

```
svcadm enable | disable system/boot-archive
```

**3 To verify the state of the boot - archive service, type:**

```
% svcs boot-archive
```

If the service is running, the output displays an online service state.

```
STATE STIME FMRI
online 9:02:38 svc:/system/boot-archive:default
```

If the service is not running, the output indicates the service is offline.

**Troubleshooting** For information about clearing automatic boot archive update failures, see [“Automatic Boot Archive Recovery” on page 270](#).

## Automatic Boot Archive Recovery

Starting with the Oracle Solaris 10 9/10 release, boot archive recovery on the SPARC platform is fully automated. On the x86 platform, boot archive recovery is partially automated.

To support auto-recovery of the boot archives on the x86 platform, a new `auto-reboot-safe` property has been added to the boot configuration SMF service, `svc:/system/boot-config:default`. By default, this property's value is set to `false`, which prevents the system from automatically rebooting to an unknown boot device. However, if your system is configured to automatically reboot to the BIOS boot device and default GRUB menu entry that the Oracle Solaris OS is installed on, you can enable automatic recovery of the boot archives by setting this property's value to `true`. The following procedure describes how to clear automatic boot archive update failures on the x86 platform.

For information about how to clear automatic boot archive update failures by using the `bootadm` command, see [“How to Clear Automatic Boot Archive Update Failures by Using the bootadm Command” on page 271](#).

### ▼ x86: How to Clear Automatic Boot Archive Update Failures by Using the `auto-reboot-safe` Property

On an x86 based systems, during the process of booting the system, if a warning similar to the following is displayed, take action as described in the following procedure.

```
WARNING: Reboot required.
The system has updated the cache of files (boot archive) that is used
during the early boot sequence. To avoid booting and running the system
with the previously out-of-sync version of these files, reboot the
system from the same device that was previously booted.
```

The system then enters system maintenance mode.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Reboot the system.**

```
reboot
```

To prevent this type of failure, if the active BIOS boot device and the GRUB menu entries point to the current boot instance, do the following:

**3 Set the auto-reboot-safe property of the svc:/system/boot-config:default SMF service to true, as follows:**

```
svccfg -s svc:/system/boot-config:default setprop config/auto-reboot-safe = true
```

**4 Verify that the auto-reboot-safe property is set correctly.**

```
svccfg -s svc:/system/boot-config:default listprop |grep config/auto-reboot-safe
config/auto-reboot-safe boolean true
```

## ▼ How to Clear Automatic Boot Archive Update Failures by Using the bootadm Command

During the process of booting the system, if a warning message that is similar to the following is displayed, take action accordingly:

```
WARNING: Automatic update of the boot archive failed.
Update the archives using 'bootadm update-archive'
command and then reboot the system from the same device that
was previously booted.
```

The following procedure describes how to manually update an out-of-date boot archive by using the bootadm command.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 To update the boot archive, type:**

```
bootadm update-archive
```

bootadm                   Manages the boot archives on a system.

update-archive           Updates the current boot archive, if required. Applies to both SPARC and x86 based systems.

**3 Reboot the system.**

```
reboot
```

## Using the bootadm Command to Manage the Boot Archives

The `/sbin/bootadm` command enables you to perform the following tasks:

- Manually update the current boot archives on a system.
- List the files and directories that are included in the boot archives on a system.
- **x86 only:** Maintain the GRUB menu.
- **x86 only:** Locate the active GRUB menu, as well as the current GRUB menu entries.

The syntax of the command is as follows:

```
/sbin/bootadm [subcommand] [-option] [-R altroot]
```

For more information about the `bootadm` command, see the [bootadm\(1M\)](#) man page.

### ▼ How to Manually Update the Boot Archive by Using the bootadm Command

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 To update the boot archive, type:**

```
bootadm update-archive
```

`bootadm`                   Manages the boot archives on a system.

`update-archive`       Updates the current boot archive, if required. Applies to both SPARC and x86 based systems.

**▪ To update the boot archive on an alternate root, type:**

```
bootadm update-archive -R /a
```

`-R altroot`       Specifies an alternate root path to apply to the `update-archive` subcommand.



---

**Note** – The root (/) file system of any non-global zone must not be referenced with the -R option. Doing so might damage the global zone's file system, compromise the security of the global zone, or damage the non-global zone's file system. See the [zones\(5\)](#) man page.

---

### 3 Reboot the system.

```
reboot
```

## ▼ How to Manually Update the Boot Archive on a Solaris Volume Manager RAID-1 (Mirrored) Root Partition

The following procedure describes how to mount a mirrored metadvice during a failsafe boot. In this procedure, the root (/) file system that is used is /dev/dsk/c0t0d0s0.

### 1 Boot the failsafe archive.

- **SPARC platform: From the ok prompt, type the following command:**

```
ok boot -F failsafe
```

If the system is already running, open a terminal window, become superuser, and type the following command:

```
reboot -- "-F failsafe"
```

For more information, see [“How to Boot a SPARC Based System in Failsafe Mode” on page 239](#).

- **x86 platform: Boot the system by selecting the failsafe boot entry in the GRUB menu.**

For more information, see [“How to Boot an x86 Based System in Failsafe Mode” on page 256](#).

Booting the system in failsafe mode produces the following output:

```
Starting shell.
#
```

### 2 During the failsafe boot, when prompted by the system to select a device to mount, type q to indicate none.

```
Please select a device to be mounted (q for none)[?,??,q]: q
```

### 3 Temporarily mount a submirror of the root (/) file system as read-only on the /a directory.

```
mount -o ro /dev/dsk/c0t0d0s0 /a
```

**4 Copy the md.conf file to the /kernel/drv directory.**

```
cp /a/kernel/drv/md.conf /kernel/drv/
```

**5 Unmount the /a directory.**

```
umount /a
```

**6 Use the devfsadm command to load the md driver.**

```
update_drv -f md
```

Running this command causes the configuration to be read, and the necessary devices to be created.

---

**Note** – Before proceeding to the next step, wait a few seconds to ensure that the md driver has had time to load.

---

**7 Use the metasync command to ensure that the root (/) file system is in sync. For example:**

```
metasync d0
```

**8 Mount the root mirror metadvice on the /a directory.**

```
mount /dev/md/dsk/d0 /a
```

**9 Update the boot archive of the device that you mounted in the previous step.**

```
bootadm update-archive -v -R /a
```

If the boot archive fails to be updated, or an error message is displayed, do the following:

**a. Update the timestamp on the md.conf file on the /a directory, which forces an update of the boot archive.**

```
touch /a/kernel/drv/md.conf
```

**b. Update the boot archive by running the bootadm command.**

```
bootadm update-archive -v -R /a
```

The update of the boot archive takes a few minutes to complete. If the boot archive was updated successfully, a message similar to the following is displayed:

```
changed /a/etc/system
cannot find: /a/etc/cluster/nodeid: No such file or directory
cannot find: /a/etc/devices/mdi_ib_cache: No such file or directory
Creating ram disk on /a
updating /a/platform/i86pc/boot_archive
```

**10 Unmount /a.**

```
umount /a
```

- 11 Reboot the system.

## ▼ How to List Contents of the Boot Archive

- 1 **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

- 2 **To list the files and directories that are included in the boot archive, type:**

```
bootadm list-archive
```

`list-archive` Lists the files and directories that are included in the boot archive or archives. Applies to both SPARC and x86 based systems.

## ▼ x86: How to Locate the Active GRUB Menu and List Current Menu Entries

Use this procedure to determine the location of the active GRUB menu and to list current GRUB menu entries.

- 1 **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

- 2 **To list the location of the active GRUB menu and current GRUB menu entries, type:**

```
bootadm list-menu
```

`list-menu` Lists the location of the active GRUB menu, as well as the current GRUB menu entries. Information about the `autoboot - timeout`, the default entry number, and the title of each entry is included in this listing. Applies to x86 based systems *only*.

### Example 13–1 Listing the Location of the Active GRUB Menu and Current GRUB Menu Entries

```
bootadm list-menu
```

```
The location for the active GRUB menu is: /stubboot/boot/grub/menu.lst
default=0
timeout=10
(0) Solaris10
(1) Solaris10 Failsafe
(2) Linux
```

## ▼ x86: How to Set the Default Boot Entry for the Active GRUB Menu

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 To set the default boot entry in the active GRUB menu, type:

```
bootadm set-menu menu-entry
```

`set-menu` Maintains the GRUB menu. The location of the active GRUB menu is `boot/grub/menu.lst`. Applies to x86 bases systems *only*.

`menu-entry` Specifies the GRUB menu entry to set as the default.

### 3 To verify default menu entry has been changed, type:

```
bootadm list-menu
```

The new default menu entry should be displayed.

#### Example 13–2 Switching the GRUB Default Menu Entry

This example shows how to switch the default GRUB menu to one of the menu entries that is displayed in the previous example. The menu entry that is selected is The Linux, menu entry 2.

```
bootadm set-menu default=2
```

**See Also** For a description of the `menu.lst` file in each GRUB implementation, see “x86: Supported GRUB Versions” on page 295.

# Troubleshooting Booting an Oracle Solaris System (Tasks)

---

This chapter describes the procedures for booting Oracle Solaris on SPARC and x86 based systems.

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

- [“Troubleshooting Booting on the SPARC Platform \(Task Map\)” on page 277](#)
- [“Troubleshooting Booting on the x86 Platform \(Task Map\)” on page 284](#)

## Troubleshooting Booting on the SPARC Platform (Task Map)

| Task                                                                                 | Description                                                                                                                                                                                                                                                        | For Instructions                                                                                                              |
|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Stop a system for recovery purposes.                                                 | If a damaged file is preventing the system from booting normally, first stop the system to attempt recovery.                                                                                                                                                       | <a href="#">“SPARC: How to Stop the System for Recovery Purposes” on page 278</a>                                             |
| Force a crash dump of and reboot of the system.                                      | You can force a crash dump and reboot of the system as a troubleshooting measure.                                                                                                                                                                                  | <a href="#">“SPARC: How to Force a Crash Dump and Reboot of the System” on page 279</a>                                       |
| Boot a SPARC based system for recovery purposes.                                     | Boot to repair an important system file that is preventing the system from booting successfully.                                                                                                                                                                   | <a href="#">“SPARC: How to Boot a System for Recovery Purposes” on page 280</a>                                               |
| Boot a SPARC based system that has an Oracle Solaris ZFS root for recovery purposes. | Boot a system to recover the root password or a similar problem that prevents you from successfully logging into an Oracle Solaris ZFS root environment, you will need to boot failsafe mode or boot from alternate media, depending on the severity of the error. | <a href="#">“SPARC: How to Boot to a ZFS Root Environment to Recover From a Lost Password or Similar Problem” on page 282</a> |
| Boot a system with the kernel debugger.                                              | You can the system with the kernel debugger to troubleshoot booting problems. Use the <code>kldb</code> command to boot the system.                                                                                                                                | <a href="#">“SPARC: How to Boot the System With the Kernel Debugger (kldb)” on page 283</a>                                   |

You might need to use one or more of the following methods to troubleshoot problems that prevent the system from booting successfully.

- Troubleshoot error messages when the system boots.
- Stop the system to attempt recovery.
- Boot a system for recovery purposes.
- Force a crash dump and reboot of the system.
- Boot the system with the kernel debugger by using the `kldb` command.

## ▼ SPARC: How to Stop the System for Recovery Purposes

### 1 Type the Stop key sequence for your system.

The monitor displays the `ok PROM` prompt.

```
ok
```

The specific Stop key sequence depends on your keyboard type. For example, you can press Stop-A or L1-A. On terminals, press the Break key.

### 2 Synchronize the file systems.

```
ok sync
```

### 3 When you see the `syncing file systems . . .` message, press the Stop key sequence again.

### 4 Type the appropriate boot command to start the boot process.

For more information, see the [boot\(1M\)](#) man page.

### 5 Verify that the system was booted to the specified run level.

```
who -r
. run-level s May 2 07:39 3 0 S
```

#### Example 14–1 SPARC: Stopping the System for Recovery Purposes

```
 Press Stop-A
ok sync
syncing file systems...
 Press Stop-A
ok boot
```

## SPARC: Forcing a Crash Dump and Reboot of the System

Forcing a crash dump and reboot of the system are sometimes necessary for troubleshooting purposes. The `savecore` feature is enabled by default.

For more information about system crash dumps, see [Chapter 17, “Managing System Crash Information \(Tasks\)”](#), in *System Administration Guide: Advanced Administration*.

## ▼ SPARC: How to Force a Crash Dump and Reboot of the System

Use this procedure to force a crash dump of the system. The example that follows this procedure shows how to use the `halt -d` command to force a crash dump of the system. You will need to manually reboot the system after running this command.

### 1 Type the stop key sequence for your system.

The specific stop key sequence depends on your keyboard type. For example, you can press Stop-A or L1-A. On terminals, press the Break key.

The PROM displays the `ok` prompt.

### 2 Synchronize the file systems and write the crash dump.

```
> n
ok sync
```

After the crash dump is written to disk, the system will continue to reboot.

### 3 Verify the system boots to run level 3.

The login prompt is displayed when the boot process has finished successfully.

```
hostname console login:
```

## Example 14–2 SPARC: Forcing a Crash Dump and Reboot of the System by Using the `halt -d` Command

This example shows how to force a crash dump and reboot of the system `jupiter` by using the `halt -d` and `boot` command. Use this method to force a crash dump and reboot of the system.

```
halt -d
Jul 21 14:13:37 jupiter halt: halted by root

panic[cpu0]/thread=30001193b20: forced crash dump initiated at user request

000002a1008f7860 genunix:kadmin+438 (b4, 0, 0, 0, 5, 0)
%l0-3: 0000000000000000 0000000000000000 0000000000000004 0000000000000004
%l4-7: 000000000000003cc 0000000000000010 0000000000000004 0000000000000004
000002a1008f7920 genunix:uadmin+110 (5, 0, 0, 6d7000, ff00, 4)
%l0-3: 0000030002216938 0000000000000000 0000000000000001 0000004237922872
%l4-7: 000000423791e770 000000000004102 0000030000449308 0000000000000005

syncing file systems... 1 1 done
dumping to /dev/dsk/c0t0d0s1, offset 107413504, content: kernel
100% done: 5339 pages dumped, compression ratio 2.68, dump succeeded
Program terminated
ok boot
Resetting ...
```

```

Sun Ultra 5/10 UPA/PCI (UltraSPARC-III 333MHz), No Keyboard
OpenBoot 3.15, 128 MB memory installed, Serial #10933339.
Ethernet address 8:0:20:a6:d4:5b, Host ID: 80a6d45b.

Rebooting with command: boot
Boot device: /pci@1f,0/pci@1,1/ide@3/disk@0,0:a
File and args: kernel/sparcv9/unix
SunOS Release 5.10 Version Generic_144500-10 64-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
configuring IPv4 interfaces: hme0.
add net default: gateway 172.20.27.248
Hostname: jupiter
The system is coming up. Please wait.
NIS domain name is example.com
.
.
.
System dump time: Wed Jul 21 14:13:41 2004
Jul 21 14:15:23 jupiter savecore: saving system crash dump
in /var/crash/jupiter/*.0
Constructing namelist /var/crash/jupiter/unix.0
Constructing corefile /var/crash/jupiter/vmcore.0
100% done: 5339 of 5339 pages saved

Starting Sun(TM) Web Console Version 2.1-dev...
.
.
.

```

## ▼ SPARC: How to Boot a System for Recovery Purposes

Use this procedure when an important file, such as `/etc/passwd`, has an invalid entry and causes the boot process to fail.

Use the stop sequence described in this procedure if you do not know the root password or if you can't log in to the system. For more information, see [“SPARC: How to Stop the System for Recovery Purposes” on page 278](#).

Substitute the device name of the file system to be repaired for the *device-name* variable in the following procedure. If you need help identifying a system's device names, refer to [“Displaying Device Configuration Information” in \*System Administration Guide: Devices and File Systems\*](#).

- 1 **Stop the system by using the system's Stop key sequence.**
- 2 **Boot the system in single-user mode.**
  - Boot the system from the Oracle Solaris installation media:
    - Insert the Oracle Solaris installation media into the drive.
    - Boot from the installation media in single-user mode.



```
ok boot cdrom -s
```

- Boot the system from the network if an installation server or remote CD or DVD drive is not available.

```
ok boot net -s
```

**3 Mount the file system that contains the file with an invalid entry.**

```
mount /dev/dsk/device-name /a
```

**4 Change to the newly mounted file system.**

```
cd /a/file-system
```

**5 Set the terminal type.**

```
TERM=sun
export TERM
```

**6 Remove the invalid entry from the file by using an editor.**

```
vi filename
```

**7 Change to the root (/) directory.**

```
cd /
```

**8 Unmount the /a directory.**

```
umount /a
```

**9 Reboot the system.**

```
init 6
```

**10 Verify that the system booted to run level 3.**

The login prompt is displayed when the boot process has finished successfully.

```
hostname console login:
```

**Example 14-3 SPARC: Booting a System for Recovery Purposes (Damaged Password File)**

The following example shows how to repair an important system file (in this case, /etc/passwd) after booting from a local CD-ROM.

```
ok boot cdrom -s
mount /dev/dsk/c0t3d0s0 /a
cd /a/etc
TERM=vt100
export TERM
vi passwd
 (Remove invalid entry)
cd /
umount /a
init 6
```

#### Example 14-4 SPARC: Booting a System If You Forgot the root Password

The following example shows how to boot the system from the network when you have forgotten the root password. This example assumes that the network boot server is already available. Be sure to apply a new root password after the system has rebooted.

```
ok boot net -s
mount /dev/dsk/c0t3d0s0 /a
cd /a/etc
TERM=vt100
export TERM
vi shadow
 (Remove root's encrypted password string)
cd /
umount /a
init 6
```

### ▼ SPARC: How to Boot to a ZFS Root Environment to Recover From a Lost Password or Similar Problem

- 1 Boot the system in failsafe mode.

```
ok boot -F failsafe
```

- 2 When prompted, mount the ZFS BE on /a.

```
.
.
ROOT/zfsBE was found on rpool.
Do you wish to have it mounted read-write on /a? [y,n,?] y
mounting rpool on /a
Starting shell.
```

- 3 Become superuser.

- 4 Change to the /a/etc directory.

```
cd /a/etc
```

- 5 Correct the passwd or shadow file.

```
vi passwd
```

- 6 Reboot the system.

```
init 6
```

## ▼ SPARC: How to Boot the System With the Kernel Debugger (kmdb)

This procedure shows you the basics for loading the kernel debugger (kmdb). For more detailed information, see the *Oracle Solaris Modular Debugger Guide*.

---

**Note** – Use the `reboot` and `halt` command with the `-d` option if you do not have time to debug the system interactively. To run the `halt` command with the `-d` option requires a manual reboot of the system afterwards. Whereas, if you use the `reboot` command, the system boots automatically. See the [reboot\(1M\)](#) for more information.

---

### 1 Halt the system, causing it to display the `ok` prompt.

To halt the system gracefully, use the `/usr/sbin/halt` command.

### 2 Type either `boot kmdb` or `boot -k` to request the loading of the kernel debugger. Press return.

### 3 Enter the kernel debugger.

The method used to enter the debugger is dependent upon the type of console that is used to access the system:

- If a locally attached keyboard is being used, press Stop-A or L1-A, depending upon the type of keyboard.
- If a serial console is being used, send a break by using the method that is appropriate for the type of serial console that is being used.

A welcome message is displayed when you enter the kernel debugger for the first time.

```
Rebooting with command: kadb
Boot device: /iommu/sbus/espdma@4,800000/esp@4,8800000/sd@3,0
.
.
.
```

#### Example 14-5 SPARC: Booting a System With the Kernel Debugger (kmdb)

```
ok boot kmdb
Resetting...

Executing last command: boot kmdb -d
Boot device: /pci@1f,0/ide@d/disk@0,0:a File and args: kmdb -d
Loading kmdb...
```

# Troubleshooting Booting on the x86 Platform (Task Map)

| Task                                            | Description                                                                                                                         | For Instructions                                                                                                     |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| Stop a system for recovery purposes.            | If a damaged file is preventing the system from booting normally, first stop the system to attempt recovery.                        | <a href="#">“x86: How to Stop a System for Recovery Purposes” on page 284</a>                                        |
| Force a crash dump of and reboot of the system. | You can force a crash dump and reboot of the system as a troubleshooting measure.                                                   | <a href="#">“x86: How to Force a Crash Dump and Reboot of the System” on page 284</a>                                |
| Boot a system with the kernel debugger.         | You can the system with the kernel debugger to troubleshoot booting problems. Use the <code>kmdb</code> command to boot the system. | <a href="#">“x86: How to Boot a System With the Kernel Debugger in the GRUB Boot Environment (kmdb)” on page 286</a> |

## ▼ x86: How to Stop a System for Recovery Purposes

### 1 Stop the system by using one of the following commands, if possible:

- If the keyboard and mouse are functional, become superuser. Then, type `init 0` to stop the system. After the `Press any key to reboot` prompt appears, press any key to reboot the system.
- If the keyboard and mouse are functional, become superuser, then, type `init 6` to reboot the system.

### 2 If the system does not respond to any input from the mouse or the keyboard, press the Reset key, if it exists, to reboot the system.

Or, you can use the power switch to reboot the system.

## x86: Forcing a Crash Dump and Reboot of the System

Forcing a crash dump and reboot of the system are sometimes necessary for troubleshooting purposes. The `savecore` feature is enabled by default.

For more information about system crash dumps, see [Chapter 17, “Managing System Crash Information \(Tasks\),” in \*System Administration Guide: Advanced Administration\*](#).

## ▼ x86: How to Force a Crash Dump and Reboot of the System

If you cannot use the `reboot -d` or the `halt -d` command, you can use the kernel debugger, `kmdb`, to force a crash dump. The kernel debugger must have been loaded, either at boot, or with the `mdb -k` command, for the following procedure to work.

---

**Note** – You must be in text mode to access the kernel debugger (kldb). So, first exit any window system.

---

### 1 Access the kernel debugger.

The method used to access the debugger is dependent upon the type of console that you are using to access the system.

- If you are using a locally attached keyboard, press F1–A.
- If you are using a serial console, send a break by using the method appropriate to that type of serial console.

The kldb prompt is displayed.

### 2 To induce a crash, use the `systemdump` macro.

```
[0]> $<systemdump
```

Panic messages are displayed, the crash dump is saved, and the system reboots.

### 3 Verify that the system has rebooted by logging in at the console login prompt.

## Example 14–6 x86: Forcing a Crash Dump and Reboot of the System by Using `halt -d`

This example shows how to force a crash dump and reboot of the x86 based system `neptune` by using the `halt -d` and `boot` commands. Use this method to force a crash dump of the system. Reboot the system afterwards manually.

```
halt -d
4ay 30 15:35:15 wacked.Central.Sun.COM halt: halted by user

panic[cpu0]/thread=ffffffff83246ec0: forced crash dump initiated at user request

fffffe80006bbd60 genunix:kadmin+4c1 ()
fffffe80006bbec0 genunix:uadmin+93 ()
fffffe80006bbf10 unix:sys_syscall32+101 ()

syncing file systems... done
dumping to /dev/dsk/clt0d0s1, offset 107675648, content: kernel
NOTICE: adpu320: bus reset
100% done: 38438 pages dumped, compression ratio 4.29, dump succeeded

Welcome to kldb
Loaded modules: [audiosup crypto ufs unix krtld s1394 sppp nca uhci lofs
genunix ip usb2 specfs nfs md random sctp]
[0]>
kldb: Do you really want to reboot? (y/n) y
```

## ▼ x86: How to Boot a System With the Kernel Debugger in the GRUB Boot Environment (kmdb)

This procedure shows the basics for loading the kernel debugger (kmdb). The `savecore` feature is enabled by default. For more detailed information about using the kernel debugger, see the [Oracle Solaris Modular Debugger Guide](#).

### 1 Boot the system.

The GRUB menu is displayed when the system is booted.

### 2 When the GRUB menu is displayed, type `e` to access the GRUB edit menu.

### 3 Use the arrow keys to select the `kernel$` line.

If you cannot use the arrow keys, use the `^` key to scroll up and the `v` key to scroll down.

### 4 Type `e` to edit the line.

The boot entry menu is displayed. In this menu, you can modify boot behavior by adding additional boot arguments to the end of the `kernel$` line.

### 5 Type `-k` at the end of the line.

### 6 Press `enter` to return to the GRUB main menu.

### 7 Type `b` to boot the system with the kernel debugger enabled.

### 8 Access the kernel debugger.

The method used to access the debugger is dependent upon the type of console that you are using to access the system:

- If you are using a locally attached keyboard, press `F1-A`.
- If you are using a serial console, send a break by using the method appropriate to that type of serial console.

A welcome message is displayed when you access the kernel debugger for the first time.

#### Example 14-7 x86: Booting a System With the Kernel Debugger (GRUB Multiboot Implementation)

This example shows how to manually boot a 64-bit capable x86 based system with the kernel debugger enabled.

```
kernel$ /platform/i86pc/multiboot kernel/amd64/unix -k -B $ZFS-BOOTFS
```

This example shows how to boot a 64-bit capable x86 based system 32-bit mode with the kernel debugger enabled.

```
kernel$ /platform/i86pc/multiboot kernel/unix -k -B $ZFS-BOOTFS
```





## x86: GRUB Based Booting (Reference)

---

This chapter contains information about x86 boot processes, including GRUB implementation details and additional GRUB reference information.

For overview information, see [Chapter 9, “Shutting Down and Booting a System \(Overview\)”](#).

For step-by-step instructions on booting a system, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\)”](#).

### x86: Boot Processes

This section includes information about boot processes that are unique to booting an x86 based system.

#### x86: System BIOS

When an x86 based system is powered on, it is controlled by the read-only-memory (ROM) Basic Input/Output System (BIOS). The BIOS is the firmware interface on Oracle Solaris operating systems that have x86 64-bit and 32-bit support.

Hardware adapters usually have an on-board BIOS that displays the physical characteristics of the device. The BIOS is used to access the device. During the startup process, the system BIOS checks for the presence of any adapter BIOS. If any adapters are found, the system then loads and executes each adapter BIOS. Each adapter's BIOS runs self-test diagnostics and then displays device information.

The BIOS on most systems has a user interface, where you can select an ordered list of boot devices that consists of the following selections:

- Diskette
- CD or DVD

- Hard disk
- Network

The BIOS attempts to boot from each device, in turn, until a valid device with a bootable program is found.

## x86: Kernel Initialization Process

The `/platform/i86pc/multiboot` program is an ELF32 executable that contains a header which is defined in the Multiboot Specification.

The multiboot program is responsible for performing the following tasks:

- Interpreting the content of boot archive
- Autodetection of systems that are 64-bit capable
- Selecting the best kernel mode for booting the system
- Assembling core kernel modules in memory
- Handing control of the system to the Solaris kernel

After the kernel gains control of the system, the kernel initializes the CPU, memory, and device subsystems. The kernel then mounts the root device, which corresponds to the `bootpath` and `fstype` properties that are specified in the `/boot/solaris/bootenv.rc` file. This file is part of the boot archive. If these properties are not specified in the `bootenv.rc` file, or on the GRUB command line, the root file system defaults to UFS on `/devices/ramdisk:a`. The root file system defaults to UFS when you boot the installation miniroot. After the root device is mounted, the kernel initializes the `sched` and `init` commands. These commands start the Service Management Facility (SMF) services.

## x86: Support for GRUB in the Oracle Solaris OS

The following sections contain additional reference information for administering GRUB in the Oracle Solaris OS.

### x86: GRUB Terminology

To thoroughly grasp GRUB concepts, an understanding of the following terms is essential.

---

**Note** – Some of the terms that are described in this list are not exclusive to GRUB based booting.

---

|                             |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>boot archive</b>         | <p>A collection of critical files that is used to boot the Oracle Solaris OS. These files are needed during system startup before the root file system is mounted. Multiple boot archives are maintained on a system:</p> <ul style="list-style-type: none"> <li>▪ A <i>primary boot archive</i> is used to boot the Oracle Solaris OS on an x86 based system.</li> <li>▪ A <i>failsafe boot archive</i> that is used for recovery when a primary boot archive is damaged. This boot archive starts the system without mounting the root file system. On the GRUB menu, this boot archive is called <i>failsafe</i>. The archive's primary purpose is to regenerate the primary boot archives, which are usually used to boot the system.</li> </ul> |
| <b>boot loader</b>          | The first software program that runs after you power on a system. This program begins the booting process.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>failsafe archive</b>     | See boot archive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>GRUB</b>                 | GNU GRand Unified Bootloader (GRUB) is an open-source boot loader with a menu interface. The menu displays a list of the operating systems that are installed on a system. GRUB enables you to easily boot these various operating systems, such as the Oracle Solaris OS, Linux, or Windows.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>GRUB main menu</b>       | A boot menu that lists the operating systems that are installed on a system. From this menu, you can easily boot an operating system without modifying the BIOS or <code>fdisk</code> partition settings.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>GRUB edit menu</b>       | A submenu of the GRUB main menu. GRUB commands are displayed on this submenu. These commands can be edited to change boot behavior.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>menu.lst file</b>        | A configuration file that lists all the operating systems that are installed on a system. The contents of this file dictate the list of operating systems that is displayed in the GRUB menu. From the GRUB menu, you can easily boot an operating system without modifying the BIOS or <code>fdisk</code> partition settings.                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>miniroot</b>             | A minimal, bootable root ( <code>/</code> ) file system that resides on the Solaris installation media. A miniroot consists of the Solaris software that is required to install and upgrade systems. On x86 based systems, the miniroot is copied to the system to be used as the failsafe boot archive. See boot archive for details about the failsafe boot archive.                                                                                                                                                                                                                                                                                                                                                                               |
| <b>primary boot archive</b> | See boot archive.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## x86: Functional Components of GRUB

GRUB consists of the following functional components:

- `stage1` – Is an image that is installed on the first sector of the `fdisk` partition. You can optionally install `stage1` on the master boot sector by specifying the `-m` option with the `installgrub` command. See the `installgrub(1M)` man page and “[Disk Management in the GRUB Boot Environment](#)” in *System Administration Guide: Devices and File Systems* for more information.
- `stage2` – Is an image that is installed in a reserved area in the `fdisk` partition. The `stage2` image is the core image of GRUB.
- `menu.lst` file – Is typically located in the `/boot/grub` directory on systems with a UFS root and in the `/pool-name/boot/grub` directory on systems with a ZFS root. This file is read by the GRUB `stage2` file. For more information, see the section, “[x86: Modifying Boot Behavior by Editing the menu.lst File](#)” on page 218.

You cannot use the `dd` command to write `stage1` and `stage2` images to disk. The `stage1` image must be able to receive information about the location of the `stage2` image that is on the disk. Use the `installgrub` command, which is the supported method for installing GRUB boot blocks.

## Naming Conventions That Are Used for Configuring GRUB

GRUB uses device-naming conventions that are slightly different from previous Solaris releases. Understanding the GRUB device-naming conventions can assist you in correctly specifying drive and partition information when you configure GRUB on your system.

The following table describes the GRUB device-naming conventions for this Oracle Solaris release.

TABLE 15-1 Conventions for GRUB Devices

| Device Name | Description                                                      |
|-------------|------------------------------------------------------------------|
| (fd0)       | First diskette                                                   |
| (fd1)       | Second diskette                                                  |
| (nd)        | Network device                                                   |
| (hd0, 0)    | First <code>fdisk</code> partition on first hard disk            |
| (hd0, 1)    | Second <code>fdisk</code> partition on first hard disk           |
| (hd0, 0, a) | Slice a on first <code>fdisk</code> partition on first hard disk |
| (hd0, 0, b) | Slice b on first <code>fdisk</code> partition on first hard disk |

---

**Note** – All GRUB device names must be enclosed in parentheses.

---

For more information about `fdisk` partitions, see “[Guidelines for Creating an fdisk Partition](#)” in *System Administration Guide: Devices and File Systems*.

## Naming Conventions That Are Used by the `findroot` Command

Starting with the Solaris 10 10/08 release, the `findroot` command replaces the `root` command that was previously used by GRUB. The `findroot` command provides enhanced capabilities for discovering a targeted disk, regardless of the boot device. The `findroot` command also supports booting from a ZFS root file system.

The following is a description of the device naming convention that is used by the `findroot` command for various GRUB implementations:

- Oracle Solaris Live Upgrade:

```
findroot (BE_x,0,a)
```

The *x* variable is the name of the boot environment.

- Standard system upgrades and new installations for systems with ZFS support:

```
findroot(pool_p,0,a)
```

The *p* variable is the name of the root pool.

- Standard system upgrades and new installations for systems with UFS support:

```
findroot (rootfsN,0,a)
```

The *N* variable is an integer number that starts at 0.

## How Multiple Operating Systems Are Supported by GRUB

This section describes how multiple operating systems that are on the same disk are supported with GRUB. The following is an example of an x86 based system that has the Solaris 10 10/08 OS, the Solaris 9 OS, Linux, and Windows installed on the same disk.

TABLE 15-2 Sample GRUB Menu Configuration

| Operating System | Location on Disk               |
|------------------|--------------------------------|
| Windows          | <code>fdisk</code> partition 0 |
| Linux            | <code>fdisk</code> partition 1 |

TABLE 15–2 Sample GRUB Menu Configuration (Continued)

| Operating System    | Location on Disk  |
|---------------------|-------------------|
| Oracle Solaris      | fdisk partition 2 |
| Solaris 9 OS        | Slice 0           |
| Solaris 10 10/08 OS | Slice 3           |

Based on the preceding information, the GRUB menu would look like the following:

```

title Oracle Solaris 10
 findroot (pool_rpool,0,a)
 kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
 module /platform/i86pc/boot_archive
title Solaris 9 OS (pre-GRUB)
 root (hd0,2,a)
 chainloader +1
 makeactive
title Linux
 root (hd0,1)
 kernel <from Linux GRUB menu...>
 initrd <from Linux GRUB menu...>
title Windows
 root (hd0,0)
 chainloader +1

```

**Note** – The Oracle Solaris slice must be the active partition. Also, do not indicate `makeactive` under the Windows menu. Doing so causes the system to boot Windows every time. Note that if Linux has installed GRUB on the master boot block, you cannot access the Oracle Solaris boot option. The inability to access the Solaris boot option occurs whether or not you designate it as the active partition.

In this case, you can do one of the following:

- Chain-load from the Linux GRUB by modifying the menu on Linux.  
*Chain-loading* is a mechanism for loading unsupported operating systems by using another boot loader.
- Replace the master boot block with the Solaris GRUB by running the `installgrub` command with the `-m` option:

```
installgrub -m /boot/grub/stage1 /boot/grub/stage2 /dev/rdisk/root-slice
```

See the `installgrub(1M)` man page for more information.

For information about the Oracle Solaris Live Upgrade boot environment, see [Oracle Solaris 10 8/11 Installation Guide: Solaris Live Upgrade and Upgrade Planning](#).

## x86: Supported GRUB Versions

In Oracle Solaris 10, GRUB uses multiboot. The contents of the `menu.lst` file vary, depending on the Oracle Solaris release you are running, the installation method that is used, and whether you are booting the system from an Oracle Solaris ZFS root or a UFS root.

- **Oracle Solaris ZFS boot support for GRUB**

If you are running a supported Oracle Solaris release, you can choose to boot from an Oracle Solaris ZFS or a UFS file system. For a description of the `menu.lst` file and an example, see “Description of the `menu.lst` File (ZFS Support)” on page 295.

- **GRUB UFS boot support**

For a description of the `menu.lst` file and an example, see “Description of the `menu.lst` File (UFS Support)” on page 296.

### Description of the `menu.lst` File (ZFS Support)

The following are various examples of a `menu.lst` file for a boot environment that contains a ZFS boot loader:

---

**Note** – Because the miniroot is mounted as the real root file system, the entry for failsafe booting in the `menu.lst` file does *not* change to the ZFS boot fs property, even if the failsafe archive is read from a ZFS dataset. The ZFS dataset is not accessed after the boot loader reads the miniroot.

---

**EXAMPLE 15-1** Default `menu.lst` File (New Installation or Standard Upgrade)

```
title Solaris 10 5/08 s10x_nbu6wos_nightly X86
findroot (pool_rpool,0,a)
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title Solaris failsafe
findroot (pool_rpool,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```

**EXAMPLE 15-2** Default `menu.lst` File (Oracle Solaris Live Upgrade)

```
title bel
findroot (BE_bel,0,a)
bootfs rpool/ROOT/szboot_0508
kernel$ /platform/i86pc/multiboot -B $ZFS-BOOTFS
module /platform/i86pc/boot_archive

title bel failsafe
findroot (BE_bel,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```

## Description of the menu.lst File (UFS Support)

The following are examples of a menu.lst file on a system that supports booting from UFS.

### EXAMPLE 15-3 Default GRUB menu.lst File (New Installation or Standard Upgrade)

```
title Solaris 10 5/08 s10x_nbu6wos_nightly X86
findroot (pool_rpool,0,a)
kernel /platform/i86pc/multiboot
module /platform/i86pc/boot_archive

title Solaris failsafe
findroot (rootfs0,0,a)
kernel /boot/multiboot kernel/unix -s -B console-ttyb
module /boot/x86.miniroot-safe
```

### EXAMPLE 15-4 Default GRUB menu.lst File (Oracle Solaris Live Upgrade)

```
title bel
findroot (BE_bel,0,a)
kernel /platform/i86pc/multiboot
module /platform/i86pc/boot_archive

title bel failsafe
findroot (BE_bel,0,a)
kernel /boot/multiboot kernel/unix -s -B console=ttyb
module /boot/x86.miniroot-safe
```



# x86: Booting a System That Does Not Implement GRUB (Tasks)

---

This chapter describes the procedures for booting an x86 based system in Oracle Solaris releases that do not implement GRUB.

---

**Note** – Starting with the Solaris 10 1/06 release, the open source GGrand Unified Bootloader (GRUB) has been implemented on x86 based systems. GRUB is responsible for loading a boot archive, which contains the kernel modules and configuration files, into the system's memory. For more information about GRUB based booting, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245.](#)

---

For overview information about the boot process, see [Chapter 9, “Shutting Down and Booting a System \(Overview\).”](#)

For step-by-step instructions on booting a SPARC based system, see [Chapter 12, “Booting an Oracle Solaris System \(Tasks\).”](#)

## x86: Booting a System (Task Map)

| Header                                        | Header                                                                                                        | Header                                                                                     |
|-----------------------------------------------|---------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Boot an x86 based system to run level 3.      | Boot to run level 3. Used after shutting down the system or performing some system hardware maintenance task. | <a href="#">“x86: How to Boot a System to Run Level 3 (Multiuser Level)” on page 299</a>   |
| Boot an x86 based system to single-user mode. | Boot to run level S. Used after performing a system maintenance task such as backing up a file system.        | <a href="#">“x86: How to Boot a System to Run Level S (Single-User Level)” on page 301</a> |

| Header                                                                                                                                                                                                                                         | Header                                                                                                                                                                                                                                                   | Header                                                                                                                                                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Boot an x86 based system interactively.                                                                                                                                                                                                        | Boot interactively. Used after making temporary changes to a system file or the kernel for testing purposes.                                                                                                                                             | <a href="#">“x86: How to Boot a System Interactively” on page 303</a>                                                                                                                                                                              |
| Boot an x86 based system from the network.                                                                                                                                                                                                     | Used to boot a PXE or non-PXE device from the network with the default network configuration strategy. This method is used for booting a diskless client.                                                                                                | <a href="#">“x86: How to Boot a System From the Network” on page 305</a>                                                                                                                                                                           |
| <p><b>Solaris 10:</b> Use the Device Configuration Assistant on an Oracle Solaris x86 based system.</p> <p><b>Note</b> – Starting with the Solaris 10 1/06 release, the Device Configuration Assistant has been replaced by the GRUB menu.</p> | Used after changing the hardware configuration of the system. This utility enables you to boot the Solaris system from a different boot device, configure new or incorrectly configured hardware, or perform other device-related or boot-related tasks. | <a href="#">“x86: How to Enter the Device Configuration Assistant” on page 308</a>                                                                                                                                                                 |
| Boot a system for recovery purposes.                                                                                                                                                                                                           | Used to boot the system when a damaged file is preventing the system from booting. You might need to do one or both of the following to boot for recovery purposes:                                                                                      | <a href="#">“x86: How to Stop a System for Recovery Purposes” on page 308</a><br><a href="#">“x86: Forcing a Crash Dump and Reboot of the System” on page 313</a><br><a href="#">“x86: How to Boot a System for Recovery Purposes” on page 308</a> |
| Boot the system with the kernel debugger (kmdb).                                                                                                                                                                                               | Used to troubleshooting system problems.                                                                                                                                                                                                                 | <a href="#">“x86: How to Boot a System With the Kernel Debugger (kmdb)” on page 311</a>                                                                                                                                                            |
| Troubleshoot boot problems on systems that have 64-bit computing capabilities.                                                                                                                                                                 | If you have hardware that requires the system to load one or more device drivers that are not available in 64-bit mode, booting the system to 64-bit mode could fail. You would then need to boot the system to 32-bit mode.                             | <a href="#">“x64: Troubleshooting a Failed 64-Bit Boot” on page 315</a>                                                                                                                                                                            |

## x86: Booting a System That Does Not Implement GRUB

The following procedures use the reset button to restart the system. If your system does not have a reset button, use the power switch to restart the system. You might be able to press Ctrl-Alt-Del to interrupt system operation, depending upon the state of the system.

## ▼ x86: How to Boot a System to Run Level 3 (Multiuser Level)

Use this procedure to boot a system that is currently at run level 0 to run level 3.

- 1 **If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

The Current Boot Parameters menu is displayed after a few minutes.

- 2 **Type `b` to boot the system to run level 3, then press Enter.**

If you do not make a selection within five seconds, the system is automatically booted to run level 3.

- 3 **Verify that the system has booted to run level 3.**

The login prompt is displayed when the boot process has finished successfully.

```
hostname console login:
```

### Example 16–1 x86: Booting a System to Run Level 3 (Multiuser Level)

For new installations, typing `b` at the boot prompt automatically boots 64-bit capable x86 based systems to 64-bit mode. For upgrades, typing `b` at the boot prompt also boots 64-bit capable x86 based systems to 64-bit mode, unless the `eeprom boot - file` parameter was previously set to a value other than `kernel/unix`.

This example shows how to boot an x86 based system that has 64-bit computing capabilities to run level 3.

```
Press any key to reboot
.
.
.
 <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:
Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults

 <<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b
SunOS Release 5.10 Version amd64-gate-2004-09-27 64-bit
Copyright 1983-2004 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
DEBUG enabled
```

```

Hostname: venus
NIS domain name is example.com
checking ufs filesystems
/dev/rdisk/c1d0s7: is logging.
venus console login:

```

### Example 16-2 x64: Manually Booting a System That Has 64-Bit Computing Capabilities in 64-Bit Mode to Run Level 3 (Multiuser Level)

For new installations, typing `b` at the boot prompt automatically boots 64-bit capable x86 based systems to 64-bit mode. For upgrades, typing `b` at the boot prompt also boots 64-bit capable x86 based systems to 64-bit mode, unless the `eeprom boot-file` parameter was previously set to a value other than `kernel/unix`.

This example shows how to *manually* boot this type of system in 64-bit mode to run level 3.

```

init 0
svc.startd: The system is coming down. Please wait.
svc.startd: 68 system services are now being stopped.
umount: /etc/svc/volatile busy
svc.startd: The system is down.
syncing file systems... done
Press any key to reboot.

Initializing system
Please wait...

 <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:

Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults

 <<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b kernel/amd64/unix
SunOS Release 5.10 Version amd64-gate-2004-09-27 64-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
DEBUG enabled
Hostname: venus
NIS domain name is example.com
checking ufs filesystems
/dev/rdisk/c1d0s7: is logging.
venus console login:

```

### Example 16-3 Manually Booting a System That Has 64-Bit Computing Capabilities in 32-Bit Mode to Run Level 3 (Multiuser Level)

For new installations, typing `b` at the boot prompt automatically boots 64-bit capable x86 based systems to 64-bit mode. For upgrades, typing `b` at the boot prompt also boots 64-bit capable x86 based systems to 64-bit mode, unless the `eeeprom boot -file` parameter was previously set to a value other than `kernel/unix`.

This example shows how to *manually* boot this type of system in 32-bit mode to run level 3.

```
init 0
svc.startd: The system is coming down. Please wait.
svc.startd: 68 system services are now being stopped.
umount: /etc/svc/volatile busy
svc.startd: The system is down.
syncing file systems... done
Press any key to reboot.
Resetting...
If the system hardware has changed, or to boot from a different
device, interrupt the autoboot process by pressing ESC.

Initializing system
Please wait...

<<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:

Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults

<<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b kernel/unix
SunOS Release 5.10 Version amd64-gate-2004-09-30 32-bit
Copyright 1983-2004 Sun Microsystems, Inc. All rights reserved.
Use is subject to license terms.
DEBUG enabled
Hostname: venus
NIS domain name is example.com
checking ufs filesystems
/dev/rdisk/c1d0s7: is logging.
venus console login:
```

## ▼ x86: How to Boot a System to Run Level S (Single-User Level)

Use this procedure to boot a system that is currently at run level 0 to run level S.

- 1 **If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

The Current Boot Parameters menu is displayed after a few minutes.

- 2 **Type `b -s` to boot the system to run level S. Press Enter.**

If you do not make a selection within five seconds, the system is automatically booted to run level 3.

- 3 **Type the superuser password, if prompted.**

- 4 **Verify that the system is at run level S.**

```
who -r
. run-level S Jul 19 14:37 S 0 3
```

- 5 **Perform the maintenance task that required the run level change to S.**

- 6 **After you complete the system maintenance task, type Control-D to bring the system to the multiuser state.**

#### Example 16-4 x86: Booting a System to Run Level S (Single-User Level)

Press any key to reboot.

Resetting...

.

Initializing system

Please wait...

<<< Current Boot Parameters >>>

Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a

Boot args:

```
Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults
```

<<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: **b -s**

SunOS Release 5.10 Version amd64-gate-2004-09-30 32-bit

Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.

Use is subject to license terms.

DEBUG enabled

Booting to milestone "milestone/single-user:default".

Hostname: venus

NIS domain name is example.com

```

Requesting System Maintenance Mode
SINGLE USER MODE

Root password for system maintenance (control-d to bypass): xxxxxx
Entering System Maintenance Mode
.
.
.
who -r
. run-level S Jul 19 14:37 S 0 3
 (Perform some maintenance task)
^D

```

## ▼ x86: How to Boot a System Interactively

Use this procedure to boot a system when you need to specify an alternate kernel or the `/etc/system` file.

- 1 **If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the reset button at this prompt. If the system is shut down, turn the system on with the power switch.

The Primary Boot Subsystem menu is displayed after a few minutes.

- 2 **Select the Solaris partition (if not marked as active) from the list. Press Enter.**

If you do not make a selection within five seconds, the active boot partition is selected automatically.

The Current Boot Parameters menu is displayed after a few minutes.

- 3 **Type `b -a` to boot the system interactively. Press Enter.**

If you do not make a selection within five seconds, the system is automatically booted to run level 3.

- 4 **Answer the following system prompts:**

- a. **When prompted, enter the name of the kernel to use for booting.**

Press enter to use the default kernel file name. Otherwise, provide the name of an alternate kernel, press Enter.

- b. **When prompted, provide an alternate path for the module directories.**

Press enter to use the default module directories. Otherwise, provide the alternate paths to module directories, press Enter.

- c. **When prompted, provide the name of an alternate system file.**  
Type `/dev/null` if your `/etc/system` file has been damaged.
  - d. **When prompted, enter the root file system type.**  
Press enter to select local disk booting with UFS, which is the default, or enter NFS for network booting.
  - e. **When prompted, enter the physical name of root device.**  
Provide an alternate device name or press return to use the default.
- 5 **If you are not prompted to answer these questions, verify that you typed the `boot -a` command correctly.**

### Example 16-5 x86: Booting a System Interactively

In the following example, the default choices (shown in square brackets [ ]) are accepted.

```
Press any key to reboot.
Resetting...
.
.
.
Autobooting from bootpath: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
```

If the system hardware has changed, or to boot from a different device, interrupt the autoboot process by pressing ESC.

```
Initializing system
Please wait...
```

```

 <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:

Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults
Running Configuration Assistant...
 <<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b -a
Enter default directory for modules [/platform/i86pc/kernel /kernel /usr/kernel]:
 Press Enter
Name of system file [etc/system]: Press Enter
SunOS Release 5.10 Version amd64-gate-2004-09-30 32-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
DEBUG enabled
```



```
root filesystem type [ufs]: Press Enter
Enter physical name of root device[/pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a]: Press
Enter

Hostname: venus
NIS domain name is example.com
checking ufs filesystems
/dev/rdisk/c1d0s7: is logging.
venus console login:
```

## x86: Booting From the Network

Any system can boot from the network if a boot server is available. You might want to boot a stand-alone system from the network for recovery purposes if the system cannot boot from the local disk.

You can boot Solaris OS x86 based systems directly from a network without the Solaris boot diskette on x86 based systems that support the Preboot Execution Environment (PXE) network booting protocol. The PXE network boot is available only for devices that implement the Intel Preboot Execution Environment specification. If the system is capable of a PXE network boot, you might want to boot the system directly from the network without using either the Device Configuration Assistant boot diskette or the Solaris Software 1 CD or DVD.

---

**Note** – In this Oracle Solaris release, the software is delivered on DVD *only*.

---

### ▼ x86: How to Boot a System From the Network

This procedure includes instructions for booting an x86 based system from the network with the Solaris Device Configuration Assistant. Note that the behavior of the Device Configuration assistant changed, starting with the Oracle Solaris OS.

Starting with the Solaris 10 1/06 release, GRUB based booting has been implemented on x86 based systems that are running the Solaris OS. The GRUB menu replaces the Device Configuration Assistant. For information about booting an x86 based system from the Network with GRUB, see [“Booting an x86 Based System From the Network” on page 260](#).

There are two network configuration strategies, Reverse Address Resolution Protocol (RARP) or Dynamic Host Configuration Protocol (DHCP). The default network boot strategy for a PXE network boot is DHCP. The default network boot strategy for non-PXE devices is RARP. For non-PXE devices, you can use either strategy, depending on whether a RARP boot server or a DHCP boot server is available on your network.

---

**Note** – If you use a DHCP server for PXE network boots, additional DHCP configuration is required. For general information on DHCP configuration, see [Part III, “DHCP,” in \*System Administration Guide: IP Services\*](#). If you want to set up your DHCP server to support installation, see [Oracle Solaris 10 8/11 \*Installation Guide: Network-Based Installations\*](#).

---

If you are performing a PXE network boot, or if you are booting the system from the Solaris Software media, the system boots automatically. The Device Configuration Assistant menu is no longer displayed by default. If you are booting a non-PXE device, you will need to follow the steps in this procedure that describe how to enter the Device Configuration Assistant menu to change the network configuration.

**1 Insert the Device Configuration Assistant boot diskette or the Solaris Software media that you want to boot from.**

▪ **Or, use the system or network adapter BIOS configuration program to enable the PXE network boot.**

- If you are using the boot diskette, the first menu of the Device Configuration Assistant is displayed.
- If you are using the Oracle Solaris Software 1 CD, DVD, or booting a PXE device from the network, the system boots automatically.

If you choose to change the network configuration and enter the Device Configuration Assistant menu, press ESC when the following message is displayed.

If the system hardware has changed, or to boot from a different device, interrupt the autoboot process by pressing ESC.

Press ESCape to interrupt autoboot in 5 seconds.

The Device Configuration Assistant screen is displayed.

**2 If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the reset button at this prompt. If the system is shut down, turn the system on with the power switch.

**3 Press the F2 key (F2\_Continue) to scan for devices.**

Device identification is performed. Then, the Identified Devices screen is displayed.

**4 Press the F2 key (F2\_Continue) to load drivers.**

Bootable drivers are loaded. Then, the Boot Solaris menu is displayed.

- 5 Use the Device Configuration Assistant to change the network configuration.
  - a. Press the F4 key (F4\_Boot Tasks).
  - b. Select Set Network Configuration Strategy. Press the F2 key (F2\_Continue).
  - c. Select either RARP or DHCP and press the F2 key (F2\_Continue).

---

**Note** – The previous step applies only if you are booting a non-PXE device from the network. For a PXE network boot, you must use DHCP, which is the default network boot strategy.

---

A message that confirms your new network boot strategy is displayed. Your network boot strategy selection is saved as the default network boot method for the next time the diskette is used for booting.

- d. Press F3\_Back to return to the Boot Solaris menu.
- 6 Select NET as the boot device.
- 7 Press F2 to continue to boot the network device.

The Solaris boot option screen is displayed.

## x86: Using the Device Configuration Assistant

---

**Note** – In this Solaris release the Device Configuration Assistant has been replaced by the GRUB menu. For more information about this feature, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245](#).

---

**Solaris 10:** The Device Configuration Assistant for the Oracle Solaris x86 based systems is a program that enables you to perform various hardware configuration and booting tasks.

You can access the Device Configuration Assistant menu from either of the following:

- Solaris boot diskette
- Oracle Solaris Software 1 CD or DVD
- PXE network boot
- Hard disk with Solaris OS installed

For the procedures in this chapter, you might be requested to insert the Device Configuration Assistant boot diskette to boot the Configuration Assistant. Alternately, if your system's BIOS supports booting from the CD or DVD, you can insert the Solaris Software 1 CD or DVD to boot the Device Configuration Assistant.

## ▼ **x86: How to Enter the Device Configuration Assistant**

**Solaris 10:** This procedure shows how to interrupt the boot process to enter the Device Configuration Assistant. In the current Solaris release, the GRUB menu replaces the Device Configuration Assistant.

### 1 **Boot the system.**

- If you are booting from the Device Configuration boot diskette, the first menu of the Device Configuration Assistant is displayed after a few minutes.
- If you are booting from the Oracle Solaris Software 1 CD, DVD, hard disk, or performing a PXE network boot, the following message is displayed:

If the system hardware has changed, or to boot from a different device, interrupt the autoboot process by pressing ESC.

Press ESCape to interrupt autoboot in 5 seconds.

If you choose to enter the Device Configuration Assistant menu, press ESC to interrupt the autoboot process.

The Device Configuration Assistant menu is displayed.

### 2 **If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the reset button at this prompt. If the system is shut down, turn the system on with the power switch.

## ▼ **x86: How to Stop a System for Recovery Purposes**

### 1 **Stop the system by using one of the following commands, if possible:**

- If the system is running, become superuser and type `init 0` to stop the system. After the **Press any key to reboot** prompt appears, press any key to reboot the system.
- If the system is running, become superuser and type `init 6` to reboot the system.

### 2 **If the system doesn't respond to any input from the mouse or keyboard, press the Reset key, if it exists, to reboot the system. Or, you can use the power switch to reboot the system.**

## ▼ **x86: How to Boot a System for Recovery Purposes**

Follow these steps to boot the system to repair a critical system resource. The example shows you how to boot from an Oracle Solaris Software CD, DVD, or from the network, mount the root (/) file system on the disk, and repair the `/etc/passwd` file.

Substitute the device name of the file system to be repaired for the *device-name* variable. If you need help identifying a system's device names, refer to “[Displaying Device Configuration Information](#)” in *System Administration Guide: Devices and File Systems*.

**1 Stop the system by using the system's Stop key sequence.**

Use the Stop key sequence for your system if you don't know the root password, or if you can't log in to the system. For more information, see “[x86: How to Stop a System for Recovery Purposes](#)” on page 308.

**2 Boot the system from the Oracle Solaris Software CD, DVD, or from the network, to single-user mode.**

**a. Insert the Device Configuration Assistant boot diskette or the Oracle Solaris Software CD or DVD that you want to boot from.**

---

**Note** – If you are using the boot diskette the Device Configuration Assistant menu is displayed. If you are using the Oracle Solaris Software CD or DVD, the system boots automatically. To enter the Device Configuration Assistant menu, press ESC to interrupt the boot process, when prompted by the system.

---

**b. If the system displays the Press any key to reboot prompt, press any key to reboot the system.**

You can also use the Reset button at this prompt. If the system is shut down, turn the system on with the power switch.

**3 The Current Boot Parameters menu is displayed after a few minutes.**

**4 Type `b -s` at the prompt. Press Enter.**

After a few minutes, the single-user mode # prompt is displayed.

**5 Mount the root (`/`) file system that contains the `invalid passwd` file.**

**6 Change to the newly mounted `etc` directory.**

**7 Make the necessary change to the file by using an editor.**

**8 Change to the root (`/`) directory.**

**9 Unmount the `/a` directory.**

**10 Reboot the system. Verify that the system has booted to run level 3.**

The login prompt is displayed when the boot process has finished successfully.

*host-name* console login:

**Example 16-6** x86: Solaris 10: Booting a System for Recovery Purposes

The following example shows how to repair the `/etc/passwd` file after booting the system automatically from a local CD-ROM in Oracle Solaris 10. GRUB based booting was introduced in the Solaris 10 1/06 release. For information about booting a system for recovery purposes in a GRUB based boot environment, see [“How to Boot an x86 Based System in Failsafe Mode” on page 256](#).

```
SunOS Secondary Boot version 3.00
```

```
Solaris Booting System
```

```
Running Configuration Assistant...
```

```
If the system hardware has changed, or to boot from a different
device, interrupt the autoboot process by pressing ESC.
```

```
Press ESCape to interrupt autoboot in 5 seconds.
```

```
Initializing system
Please wait...
```

```
<<< Current Boot Parameters >>>
```

```
Boot path: /pci@0,0/pci-ide@7,1/ide@1/sd@0,0:a
```

```
Boot args:
```

```
Select the type of installation you want to perform:
```

- 1 Solaris Interactive
- 2 Custom JumpStart
- 3 Solaris Interactive Text (Desktop session)
- 4 Solaris Interactive Text (Console session)

```
Enter the number of your choice followed by the <ENTER> key.
Alternatively, enter custom boot arguments directly.
```

```
If you wait for 30 seconds without typing anything,
an interactive installation will be started.
```

```
Select type of installation: b -s
```

```
.
.
.
```

```
mount /dev/dsk/c0t0d0s0 /a
```

```

.
.
.
cd /a/etc
vi passwd
 (Remove invalid entry)
cd /
umount /a
init 6

```

## ▼ x86: How to Boot a System With the Kernel Debugger (kldb)

This procedure shows the basics for loading the kernel debugger (kldb) in Oracle Solaris 10. The `savecore` feature is enabled by default. For more detailed information about using the kernel debugger, see the *Oracle Solaris Modular Debugger Guide*.

For step-by-step instructions on booting a system with the kernel debugger in the current Solaris release, see “x86: How to Boot a System With the Kernel Debugger in the GRUB Boot Environment (kldb)” on page 286.

- 1 **Boot the system.**
- 2 **Type `b -k` at the `Select (b)oot or (i)nterpreter` prompt. Press `Enter`.**
- 3 **Access the kernel debugger.**

The method used to enter the debugger is dependent upon the type of console that is used to access the system:

- If a locally attached keyboard is being used, press `F1-A`.
- If a serial console is being used, send a break by using the method appropriate to the type of serial console that is being used.

A welcome message is displayed when you access the kernel debugger for the first time.

### Example 16–7 x86: Booting a System With the Kernel Debugger (kldb)

Typing `b -k` at the `Select (b)oot or (i)nterpreter` boot prompt boots a system to its default mode and also loads `kldb`. This example shows how to boot an x86 based system that has 32-bit computing capabilities to 32-bit mode and also load `kldb`.

Press any key to reboot.

```

.
.
.
<<< Current Boot Parameters >>>

```

```

Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:

Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults
Running Configuration Assistant...
 <<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b -k
Loading kmdb...
SunOS Release 5.10 Version gate:2004-10-21 32-bit
Copyright (c) 1983, 2011, Oracle and/or its affiliates. All rights reserved.
Use is subject to license terms.
.
.
.

```

### Example 16-8 x64: Manually Booting a System That Has 64-Bit Computing Capabilities to 64-Bit Mode With the Kernel Debugger (kmdb)

This example shows how to manually boot an x86 based system that has 64-bit computing capabilities to 64-bit mode with kmdb.

Press any key to reboot

```

.
.
.
 <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:
Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults

 <<< timeout in 5 seconds >>>

Select (b)oot or (i)nterpreter: b kernel/amd64/unix -k
Loading kmdb...

```

### Example 16-9 32-bit x64: Manually Booting a System That Has 64-Bit Computing Capabilities to 32-Bit Mode With the Kernel Debugger (kmdb)

This example shows how to manually boot an x86 based system that has 64-bit computing capabilities to 32-bit mode with kmdb.

Press any key to reboot

```

.
.
.
 <<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:

```



|      |                                    |                           |
|------|------------------------------------|---------------------------|
| Type | b [file-name] [boot-flags] <ENTER> | to boot with options      |
| or   | i <ENTER>                          | to enter boot interpreter |
| or   | <ENTER>                            | to boot with defaults     |

```
<<< timeout in 5 seconds >>>
```

```
Select (b)oot or (i)nterpreter: b kernel/unix -k
Loading kmdb...
```

## x86: Forcing a Crash Dump and Reboot of the System

Forcing a crash dump and rebooting the system is sometimes necessary for troubleshooting purposes. The `savecore` feature is enabled by default.

For more information on system crash dumps, see [Chapter 17, “Managing System Crash Information \(Tasks\)”](#) in *System Administration Guide: Advanced Administration*.

### ▼ x86: How to Force a Crash Dump and Reboot of the System

If you cannot use the `reboot -d` or the `halt -d` command, you can use the kernel debugger, `kmdb`, to force a crash dump. The kernel debugger must have been loaded, either at boot, or with the `mdb -k` command, for the following procedure to work.

---

**Note** – You must be in text mode to enter the kernel debugger (`kmdb`). So, first exit any window system.

---

- 1 If a locally-attached keyboard is being used as the system console, press F1-A on that keyboard. If the system is configured to use a remote (serial) console, use the mechanism that is appropriate to that console to send a break character.**

The `kmdb` prompt is displayed.

- 2 Use the `systemdump` macro to induce a crash.**

```
[0]> $<systemdump
```

Panic messages are displayed, the crash dump is saved, and the system reboots.

- 3 Verify that the system has rebooted by logging in at the console login prompt.**

#### Example 16–10 x86: Forcing a Crash Dump and Reboot of the System by Using `halt -d`

This example shows how to force a crash dump and reboot of the x86 based system `neptune` by using the `halt -d` and `boot` commands. Use this method to force a crash dump of the system. You will need to manually reboot the system after running the `halt` command with the `-d` option.

```
halt -d
```

```
Aug 11 12:51:27 neptune halt:
```

```
halted by <user> panic[cpu45]/thread=d3971a00: forced crash dump initiated at user request
```

```
d363ae58 genunix:kadmin+bd (5, 0, 0, d3fefac0)
```

```
d363af88 genunix:uadmin+88 (5, 0, 0, 0, 0, d363afb4)
```

```
syncing file systems... done
```

```
dumping to /dev/dsk/c0t0d0s1, offset 107806720, content: kernel
```

```
100% done: 40223 pages dumped, compression ratio 4.11, dump succeeded
```

```
Press any key to reboot.
```

```
Resetting...
```

```
.
```

```
.
```

```
SunOS Secondary Boot version 3.00
```

```
Autobooting from bootpath: /pci@0,0/pci1028,10a@3/sd@0,0:a
```

```
Running Configuration Assistant..
```

```
If the system hardware has changed, or to boot from a different
device, interrupt the autoboot process by pressing ESC.
```

```
Initializing system
```

```
Please wait...
```

```
<<< Current Boot Parameters >>>
```

```
Boot path: /pci@0,0/pci1028,10a@3/sd@0,0:a
```

```
Boot args:
```

```
Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults
```

```
<<< timeout in 5 seconds >>>
```

```
Select (b)oot or (i)nterpreter:
```

```
Loading kmdb...
```

```
SunOS Release 5.10 Version s10_62 32-bit
```

```
Copyright 1983-2004 Sun Microsystems, Inc. All rights reserved.
```

```
Use is subject to license terms.
```

```
configuring IPv4 interfaces: iprb0.
```

```
add net default: gateway 172.20.26.248
```

```
Hostname: neptune
```

```
The system is coming up. Please wait.
```

```
checking ufs filesystems
```

```
/dev/rdisk/c0t0d0s7: is logging.
```

```
NIS domain name is example.com
```

```
starting rpc services: rpcbind keyserv ypbind done.
```

```
Setting netmask of iprb0 to 255.255.255.0
```

```
Setting default IPv4 interface for multicast: add net 224.0/4: gateway venus
```

```
syslog service starting.
```

```
System dump time: Wed Aug 11 12:51:29 2004
```

```
Aug 11 13:13:26 venus savecore: saving system crash dump in /var/crash/venus/*.1
```

```
Constructing namelist /var/crash/venus/unix.1
```

```
Constructing corefile /var/crash/venus/vmcore.1
```

```
100% done: 42157 of 42157 pages saved
```

```
volume management starting.
```

```
The system is ready.
```

```
.
```

## x64: Troubleshooting a Failed 64-Bit Boot

In some instances, an attempt to boot a 64-bit capable x86 based system to 64-bit mode might fail. This failure might produce an error similar to the following:

```
Select (b)oot or (i)nterpreter: b kernel/amd64/unix
.
.
.
pci: cannot load driver
Cannot load drivers for /pci@0,0/pci1022,7450@a/pci17c2,10@4/sd@0,0:a
(Can't load the root filesystem)
Press any key to reboot.
.
.
.
```

In the event such a failure occurs, boot the system to 32-bit mode by typing the following command at the Select (b)oot or (i)nterpreter boot prompt:

```
Select (b)oot or (i)nterpreter: b kernel/unix
```

For more information, see [Example 16-3](#).

## x86: Boot Processes (Reference)

The following sections include reference information that pertains to booting a Solaris x86 based system that does *not* implement GRUB based booting.

---

**Note** – The GRUB menu has replaced the Solaris Device Configuration Assistant in this release. For more information about booting an x86 based system in this Oracle Solaris release, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245](#).

---

## x86: Boot Subsystems

During the boot process, the boot subsystem menus allow you to customize boot choices. If the system receives no response during the timeout periods, it continues to boot automatically using the default selections. You can stop the boot process when each boot subsystem menu is displayed. Or, you can let the boot process continue automatically.

At three points during the Solaris boot process, you can make the following choices about a booting system:

- **Primary Boot Subsystem (Partition Boot Menu)** – This first menu appears if multiple operating systems exist on the disk. The menu enables you to boot any of the operating systems installed. By default, the operating system that is designed as *active* is booted.
 

Note that if you choose to boot a system other than the Oracle Solaris OS, you cannot reach the next two menus.
- **Interrupt the Autoboot Process** – If the autoboot process is interrupted, you can access the Device Configuration Assistant menu.
 

The Solaris Device Configuration Assistant enables you to boot the Solaris system from a different boot device, configure new or misconfigured hardware, or perform other device-related or boot-related tasks.
- **Current Boot Parameters menu** – Two forms of this menu exist, one menu for a normal Solaris boot and one menu for a Solaris installation boot:
  - The normal Current Boot Parameters menu enables you to boot the Solaris system with options, or enter the boot interpreter.
  - The install Current Boot Parameters menu enables you to select the type of installation to be performed or to customize the boot process.

The following table summarizes the purpose of the primary x86 based system boot interfaces. See the sections that follow for a detailed description and example of each boot interface.

TABLE 16-1 x86: Boot Subsystems

| Boot Subsystem                                       | Purpose                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Primary Boot Subsystem (Partition Boot menu)         | This menu appears if the disk you are booting from contains multiple operating systems, including the Oracle Solaris OS (Solaris OS).                                                                                                                                                                                                                                                       |
| Secondary Boot Subsystem                             | This menu appears each time you boot the Oracle Solaris release. The system is booted automatically, unless you choose to run the Solaris Device Configuration Assistant by interrupting the autoboot process.                                                                                                                                                                              |
| Solaris Device Configuration Assistant/Boot Diskette | There are two ways to access the Device Configuration Assistant menus: <ul style="list-style-type: none"> <li>■ Use the Device Configuration Assistant boot diskette or the Oracle Solaris Software CD (on systems that can boot from the CD-ROM drive) to boot the system.</li> <li>■ Interrupt the autoboot process when you boot the Solaris software from an installed disk.</li> </ul> |
| Current Boot Parameters menu                         | This menu appears when you boot the Oracle Solaris release from the software media, or the network. The menu presents a list of boot options.                                                                                                                                                                                                                                               |

---

**Note** – If you need to create the Solaris Device Configuration Assistant boot diskette, go to <http://www.oracle.com/webfolder/technetwork/hcl/index.html>.

---

## x86: Booting the Solaris Release

In this release, if you are booting an x86 based system with the Oracle Solaris Software CD, DVD, or performing a PXE network boot, the system will boot automatically. To use the Device Configuration Assistant, you must interrupt the boot process by pressing ESC when prompted by the system.

During the device identification phase, the Device Configuration Assistant does the following:

- Scans for devices that are installed on the system
- Displays the identified devices
- Enables you to perform optional tasks such as selecting a keyboard type or editing devices and their resources

During the boot phase, the Device Configuration Assistant does the following:

- Displays a list of devices from which to boot. The device marked with an asterisk (\*) is the default boot device.
- Enables you to perform optional tasks, such as editing autoboot settings and property settings, and choosing the network configuration strategy.

The following section provides examples of menus that appear during the device identification phase. The device output varies based on your system configuration.

## x86: Screens Displayed During the Device Identification Phase

Several screens are displayed as the Device Configuration Assistant attempts to identify devices on the system.

This section provides examples of the following boot subsystem screens:

- Device Configuration Assistant screen
- Bus Enumeration screen
- Scanning Devices screen
- Identified Devices screen

---

## x86: Device Configuration Assistant Screen

---

**Note** – In this Oracle Solaris release, the Device Configuration Assistant Screen has been replaced with the GRUB menu on x86 based systems. For more information, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245](#).

---

In the Solaris 10 initial release, the autoboot process bypasses the Device Configuration Assistant menus, unless you press ESC when prompted by the system during the boot phase. If you choose to use the Device Configuration Assistant, the following screen is displayed.

```
Solaris Device Configuration Assistant
```

```
The Solaris(TM)Device Configuration Assistant
scans to identify system hardware, lists identified devices, and can
boot the Solaris software from a specified device. This program must be
used to install the Solaris operating environment, add a driver,
or change the hardware on the system.
```

```
> To perform a full scan to identify all system hardware, choose Continue.
> To diagnose possible full scan failures, choose Specific Scan.
> To add new or updated device drivers, choose Add Driver.
```

```
About navigation...
```

- The mouse cannot be used.
- If the keyboard does not have function keys or they do not respond, press ESC. The legend at the bottom of the screen will change to show the ESC keys to use for navigation.
- The F2 key performs the default action.

```
F2_Continue F3_Specific Scan F4_Add Driver F6_Help
```

## x86: Bus Enumeration Screen

The Bus Enumeration screen appears briefly while the Device Configuration Assistant gathers hardware configuration data for devices that can be detected automatically.

```
Bus Enumeration
```

```
Determining bus types and gathering hardware configuration data ...
```

```
Please wait ...
```

## x86: Scanning Devices Screen

The Scanning Devices screen appears while the Device Configuration Assistant manually scans for devices that can only be detected with special drivers.

```
Scanning Devices
```

```
The system is being scanned to identify system hardware.
```

If the scanning stalls, press the system's reset button. When the system reboots, choose Specific Scan or Help.

Scanning: Floppy disk controller

```
#####
| | | | |
0 20 40 60 80 100
```

Please wait ...

## x86: Identified Devices Screen

The Identified Devices screen displays which devices have been identified on the system. From here, you can continue to the Boot Solaris menu.

Or, you can perform the following optional device tasks:

- Setting a keyboard configuration
- Viewing and editing devices
- Setting up a serial console
- Saving and deleting configurations

### Identified Devices

The following devices have been identified on this system. To identify devices not on this list or to modify device characteristics, such as keyboard configuration, choose Device Tasks. Platform types may be included in this list.

```
ISA: Floppy disk controller
ISA: Motherboard
ISA: PnP bios: 16550-compatible serial controller
ISA: PnP bios: 16550-compatible serial controller
ISA: PnP bios: Mouse controller
ISA: PnP bios: Parallel port
ISA: System keyboard (US-English)
PCI: Bus Mastering IDE controller
PCI: Universal Serial Bus
PCI: VGA compatible display adapter
```

F2\_Continue F3\_Back F4\_Device Tasks F6\_Help

## x86: Menus Displayed During the Boot Phase

---

**Note** – Starting with the Solaris 10 1/06 release the GRUB is displayed when the system is booted. For more information about GRUB based booting, see [“Booting an x86 Based System by Using GRUB \(Task Map\)” on page 245](#).

---

During this phase, you can determine the way in which the system is booted.

The following menus are displayed during the boot phase:

- Boot Solaris menu
- Current Boot Parameters menu

## x86: Boot Solaris Menu

The Boot Solaris menu allows you to select the device from which to boot the Oracle Solaris release. You can also perform optional tasks, such as viewing and editing autoboot and property settings. Once you select a boot device and you choose Continue, the Solaris kernel begins to boot.

### Boot Solaris

Select one of the identified devices to boot the Solaris kernel and choose Continue.

To perform optional features, such as modifying the autoboot and property settings, choose Boot Tasks.

An asterisk (\*) indicates the current default boot device.

> To make a selection use the arrow keys, and press Enter to mark it [X].

```
[X] DISK: (*) Target 0:QUANTUM FIREBALL1280A
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
[] DISK: Target 1:ST5660A
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
[] DISK: Target 0:Maxtor 9 0680D4
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
[] CD : Target 1:TOSHIBA CD-ROM XM-5602B 1546
on Bus Mastering IDE controller on Board PCI at Dev 7, Func 1
```

```
F2_Continue F3_Back F4_Boot Tasks F6_Help
```

## x86: Current Boot Parameters Menu

This menu appears each time you boot the Oracle Solaris release from the local disk. Let the five-second timeout elapse if you want to boot the default kernel. If you want to boot with different options, select an appropriate option before the timeout period elapses.

```
<<< Current Boot Parameters >>>
Boot path: /pci@0,0/pci-ide@7,1/ide@0/cmdk@0,0:a
Boot args:
Type b [file-name] [boot-flags] <ENTER> to boot with options
or i <ENTER> to enter boot interpreter
or <ENTER> to boot with defaults
```

```
<<< timeout in 5 seconds >>>
```

Select (b)oot or (i)nterpreter:



## x86: Boot Process

The following table describes the boot process on x86 based systems.

TABLE 16–2 x86: Description of the Boot Process

| Boot Phase            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BIOS                  | <p>1. When the system is turned on, the BIOS runs self-test diagnostics to verify the system's hardware and memory. The system begins to boot automatically if no errors are found. If errors are found, error messages are displayed that describe recovery options.</p> <p>The BIOS of additional hardware devices are run at this time.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Boot Programs         | <p>2. The BIOS boot program tries to read the first disk sector from the boot device. This first disk sector on the boot device contains the master boot record <code>mboot</code>, which is loaded and executed. If no <code>mboot</code> file is found, an error message is displayed.</p> <p>3. The master boot record, <code>mboot</code>, contains disk information needed to find the active partition and the location of the Solaris boot program, <code>pboot</code>, loads and executes <code>pboot</code>, <code>mboot</code>.</p> <p>4. The Solaris boot program, <code>pboot</code>, loads <code>bootblk</code>, the primary boot program. The purpose of <code>bootblk</code> is to load the secondary boot program, which is located in the UFS file system.</p> <p>5. If there is more than one bootable partition, <code>bootblk</code> reads the <code>fdisk</code> table to locate the default boot partition, and builds and displays a menu of available partitions. You have a 30 seconds to select an alternate partition from which to boot. This step occurs only if there is more than one bootable partition present on the system.</p> <p>6. <code>bootblk</code> finds and executes the secondary boot program, <code>boot.bin</code> or <code>ufsboot</code>, in the root (<code>/</code>) file system. You have five seconds to interrupt the autoboot to start the Solaris Device Configuration Assistant.</p> <p>7. The secondary boot program, <code>boot.bin</code> or <code>ufsboot</code>, starts a command interpreter that executes the <code>/etc/bootrc</code> script. This script provides a menu of choices for booting the system. The default action is to load and execute the kernel. You have a 5-second interval to specify a boot option or to start the boot interpreter.</p> |
| Kernel initialization | <p>8. The kernel initializes itself and begins loading modules by using the secondary boot program (<code>boot.bin</code> or <code>ufsboot</code>) to read the files. When the kernel has loaded enough modules to mount the root (<code>/</code>) file system, the kernel unmaps the secondary boot program and continues, using its own resources.</p> <p>9. The kernel creates a user process and starts the <code>/sbin/init</code> process, which starts other processes by reading the <code>/etc/inittab</code> file.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

TABLE 16–2 x86: Description of the Boot Process (Continued)

| Boot Phase | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| init       | <p>10. In this Oracle Solaris release, the <code>/sbin/init</code> process starts <code>/lib/svc/bin/svc.startd</code>, which starts system services that do the following:</p> <ul style="list-style-type: none"> <li>▪ Check and mount file systems</li> <li>▪ Configure network and devices</li> <li>▪ Start various processes and perform system maintenance tasks</li> </ul> <p>In addition, <code>svc.startd</code> executes the run control (<code>rc</code>) scripts for compatibility.</p> |

## x86: Boot Files

In addition to the run control scripts and boot files, there are additional boot files that are associated with booting x86 based systems.

TABLE 16–3 x86: Boot Files

| File                                    | Description                                                                                                                    |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| <code>/etc/bootrc</code>                | Contains menus and options for booting the Oracle Solaris release.                                                             |
| <code>/boot</code>                      | Contains files and directories needed to boot the system.                                                                      |
| <code>/boot/mdboot</code>               | DOS executable that loads the first-level bootstrap program ( <code>strap.com</code> ) into memory from disk.                  |
| <code>/boot/mdbootbp</code>             | DOS executable that loads the first-level bootstrap program ( <code>strap.com</code> ) into memory from diskette.              |
| <code>/boot/rc.d</code>                 | Directory that contains install scripts. Do not modify the contents of this directory.                                         |
| <code>/boot/solaris</code>              | Directory that contains items for the boot subsystem.                                                                          |
| <code>/boot/solaris/boot.bin</code>     | Loads the Solaris kernel or stand-alone <code>kmdb</code> . In addition, this executable provides some boot firmware services. |
| <code>/boot/solaris/boot.rc</code>      | Prints the Oracle Solaris Operating OS on an x86 system and runs the Device Configuration Assistant in DOS-emulation mode.     |
| <code>/boot/solaris/bootconf.exe</code> | DOS executable for the Device Configuration Assistant.                                                                         |
| <code>/boot/solaris/bootconf.txt</code> | Text file that contains internationalized messages for Device Configuration Assistant ( <code>bootconf.exe</code> ).           |
| <code>/boot/solaris/bootenv.rc</code>   | Stores eeprom variables that are used to set up the boot environment.                                                          |

TABLE 16-3 x86: Boot Files (Continued)

---

| File                    | Description                                                                                                  |
|-------------------------|--------------------------------------------------------------------------------------------------------------|
| /boot/solaris/devicedb  | Directory that contains the master file, a database of all possible devices supported with realmode drivers. |
| /boot/solaris/drivers   | Directory that contains realmode drivers.                                                                    |
| /boot/solaris/itup2.exe | DOS executable run during install time update (ITU) process.                                                 |
| /boot/solaris/machines  | Obsolete directory.                                                                                          |
| /boot/solaris/nbp       | File associated with network booting.                                                                        |
| /boot/solaris/strap.rc  | File that contains instructions on what load module to load and where in memory it should be loaded.         |
| /boot/strap.com         | DOS executable that loads the second-level bootstrap program into memory.                                    |

---



## Working With the Oracle Solaris Auto Registration `regadm` Command (Tasks)

---

This chapter describes how to use the `regadm` command, which is part of the new Oracle Solaris Auto Registration feature as of Oracle Solaris 10 9/10.

Auto Registration is a built-in mechanism that enables configuration data about your system to be gathered and automatically registered with the Oracle Product Registration System for inventory and support purposes. The `regadm` command-line interface (CLI) can be used by a privileged system administrator to manage the Auto Registration feature and to explicitly register a system, independent of the installation or upgrade process.

---

**Note** – This chapter does *not* provide detailed information about auto registration processes that are part of an installation or upgrade. Nor does the chapter contain information about which installation methods and configurations support Auto Registration.

For overview information about Auto Registration and details about how the feature works with the installation process, see [“Oracle Solaris Auto Registration”](#) in *Oracle Solaris 10 8/11 Installation Guide: Planning for Installation and Upgrade*.

---

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

- “Oracle Solaris Auto Registration (Overview)” on page 326
- “Getting Started With the `regadm` Command” on page 330
- “Managing the Auto Registration SMF Service” on page 331
- “Configuring the Auto Registration Feature by Using the `regadm` Command” on page 333
- “Registering Your System by Using the `regadm` Command” on page 341
- “Troubleshooting Auto Registration Issues” on page 345

## Oracle Solaris Auto Registration (Overview)

Oracle Solaris Auto Registration is a mechanism that is built into Oracle Solaris 10 9/10. This mechanism enables basic asset information about your system, called a *service tag*, to be transmitted to the Oracle Product Registration System upon reboot after an installation or an upgrade, and again on subsequent reboots, if any system configuration changes have occurred. The Oracle Product Registration System consists of web services that are deployed on the public Internet. These services are responsible for implementing the registration of a system's assets. Oracle uses this information to improve its products and services. You can also retrieve and use this information to more effectively manage your inventory.

Auto Registration leverages the existing service tag technology that enables a system's assets to be discovered on a network and then registered in a local registry. Service tags contain information about the products that are installed on your system, for example, the OS version and the hardware configuration. Each service tag describes a different product. To enable the sending of service tags, use your My Oracle Support credentials. For more information, see “[How Your Oracle Solaris 10 System Is Registered](#)” on page 326.

For more information about how service tags work, go to <http://www.oracle.com/technetwork/server-storage/solaris/oracle-service-tag-faq-418684.html>.

For general information about service tags, see the `stclient(1M)` man page.

For answers to common questions and information about how to gain initial access to My Oracle Support, go to <http://www.oracle.com/us/support/044753.html#SS6>.

## How Your Oracle Solaris 10 System Is Registered

Oracle uses support credentials and network connectivity information that is collected before or during an installation or upgrade to associate product registration with a specific user account. Support credentials can also be provided after an installation by using the `regadm auth` command. See “[Authenticating a User by Specifying Named Support Credentials](#)” on page 338. If no support credentials or network connectivity information is specified, the Auto Registration process assumes a direct HTTP network connection (no proxy), and the registration is anonymous.

To register your system by specifying support credentials, you must have either a valid My Oracle Support account or a valid Sun Online account. Support credentials consist of a user name and password that enable you to authenticate to Oracle's support portal and the Oracle Product Registration System.

For information about obtaining a My Oracle Support account, go to <http://www.oracle.com/support/premier/myoraclesupport.html>.

For information about obtaining a Sun Online account, go to <https://reg.sun.com/register>.

If you are running Oracle Solaris 10 9/10, you can provide support credentials, as follows:

- **Before or during an installation or upgrade**

- Before an automatic installation or upgrade (by using either a network installation or Oracle Solaris JumpStart), you can provide support credentials by adding the new `auto_reg` keyword to your `sysidcfg` file. The same keyword can be used to set up an anonymous registration or to disable the Auto Registration feature. If you do not add the `auto_reg` keyword to the `sysidcfg` file, you are prompted to provide your support credentials during the installation or upgrade process.
- During an interactive installation or upgrade, you are prompted for your support credentials. Or, you can register anonymously. During an interactive installation or upgrade, you can also provide network configuration information, such as the proxy server and proxy port.
- During an installation or upgrade, you can use the new Oracle Solaris Live Upgrade `-k` command option to provide your support credentials and proxy information, or to register anonymously.

For more information about providing support credentials before or during an installation, see [“Oracle Solaris Auto Registration”](#) in *Oracle Solaris 10 8/11 Installation Guide: Planning for Installation and Upgrade*.

- **After an installation or upgrade**

After an installation, you can use the `regadm` command to transition a system from being anonymously registered to fully-registered by providing named support credentials. You can also use the `regadm` command to provide network configuration information, such as proxy server and proxy port. For more information, see [“Authenticating a User by Specifying Named Support Credentials”](#) on page 338.

## When and How Service Tag Data Is Transmitted

Auto Registration is managed by the `svc:/application/autoreg` SMF service. This service runs once at boot time to check for changes in your system's configuration.

If any changes have occurred, the service tag data is automatically sent to the Oracle Product Registration System, as follows:

- **During an initial reboot after an installation or upgrade**

When you reboot your system after an installation or upgrade, information about newly installed products is automatically sent to the Oracle Product Registration System.

- **During subsequent reboots when system configuration changes have occurred**

After an installation or upgrade, if the product database has changed, information is automatically sent to the Oracle Product Registration System.

- **Explicitly, by using the `regadm register` command**

Using the `regadm register` command to register your system is essentially the same as performing a reboot of the system. If any system configuration changes have occurred, product registration occurs immediately upon running the command. The configuration data is transmitted to the Oracle Product Registration System automatically, without needing to reboot your system. For more information about using the `regadm` command to register your system, see “[How to Register Your System by Using the regadm Command](#)” on page 342.

To obtain information about your registered products, go to <https://inventory.sun.com/inventory/data.jsp>. For general information about service tags, go to <http://www.oracle.com/technetwork/server-storage/solaris/oracle-service-tag-faq-418684.html>.

## Managing Auto Registration Configuration

After an installation or upgrade, Auto Registration configuration can be managed by using the `regadm` command.

After an installation, a privileged system administrator can perform the following tasks by using the `regadm` command:

- Administer the SMF service that manages the Auto Registration process.
- Display the current Auto Registration configuration.
- Authenticate a specific user or account with My Oracle Support by specifying named support credentials.
- Configure or reconfigure Auto Registration properties.
- Clear Auto Registration properties.
- Explicitly register your system, independent of the installation process.

For more information, see the [regadm\(1M\)](#) man page.

For task-related information, see “[Registering Your System by Using the regadm Command](#)” on page 341.

## How the Auto Registration Process Works

During a system reboot, the Auto Registration SMF service checks your system's configuration against the information in the Service Tag Registry to determine whether any changes have occurred since the last reboot. If changes are discovered, a service tag for the newly installed product, for example, the Oracle Solaris 10 9/10 OS, is sent to the Oracle Product Registration System and registered with your recorded support credentials, or anonymously, if no support



credentials were provided. The data is transmitted by using a Hypertext Transfer Protocol Secure (HTTPS) connection. If no changes are detected during a reboot, no registration is performed.

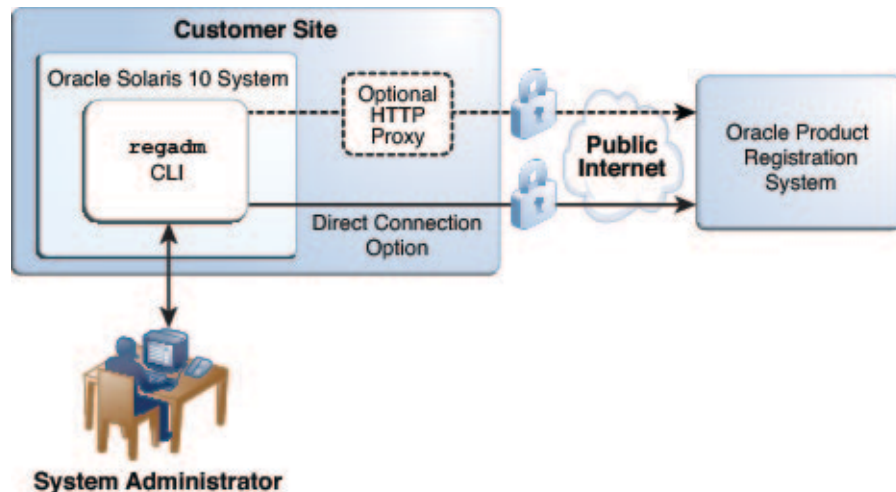
Although Auto Registration feature is enabled by default, you can optionally disable the feature. For information, see [“How to Disable the Auto Registration SMF Service”](#) on page 332.

---

**Note** – The Auto Registration SMF service runs in global zones *only*.

---

In addition to the Auto Registration SMF service, the Auto Registration process consists of several other components that work together to affect auto registration of your system's assets. The following figure illustrates these components and their relationship to each other.



Auto Registration consists of the following components:

- **Auto Registration command-line interface**

The Auto Registration CLI (`/usr/sbin/regadm`) enables a privileged system administrator Auto Registration on an Oracle Solaris system. Single commands can be executed from the command line or automated through the use of a script in large server environments. For more information, see the `regadm(1M)` man page and [“Registering Your System by Using the regadm Command”](#) on page 341.

- **Oracle Product Registration System**

The Oracle Product Registration System consists of web services that are deployed on the public Internet. These services are responsible for implementing the registration of products.

- **Secure transport mechanism**

Oracle uses a secure, one-way transport system for access to its registration system. Client-generated information is encrypted with a private key. Registered products are extracted from the Service Tag Registry and then uploaded to My Oracle Support through the use of a secure HTTPS connection (port 443). The use of an HTTPS connection creates a secure channel over an insecure network. If no HTTP proxy and port are specified during registration, a direct network connection is used. Communications are outbound and are initiated from a customer system only.

- **Auto Registration data elements**

The data elements that comprise a registration are identical to the data elements that are used in the Service Tag schema. For more information about the data that is transmitted to the Oracle Product Registration System, go to <https://inventory.sun.com/inventory/>.

## Getting Started With the `regadm` Command

The following is information that will help you get started with administering Auto Registration by using the `regadm` command.

### When to Use the `regadm` Command-Line Interface

The following are situations in which you might need to use the `regadm` command:

- **To enable the Auto Registration feature**

In this situation, you previously disabled Auto Registration, but you now want to use it. To do so, you must enable the Auto Registration SMF service by using the `regadm enable` command.

- **To attach named support credentials to a registered system for support purposes**

You previously registered products anonymously, but for support purposes, you now need to attach named support credentials to these products.

- **To explicitly register a system, independent of the installation or upgrade process**

You want to explicitly register your system, independent of the installation or upgrade process.

---

**Note** – When you use the `regadm register` command to register your system, the service tag data is transmitted to the Oracle Product Registration System immediately upon running the command. There is no need to reboot the system for a service tag to be sent.

---

## Requirements for Using the `regadm` Command

The following are the requirements for using the `regadm` command:

- You must become superuser or assume an equivalent role.  
The `regadm` command is also assigned to the Software Installation rights profile.
- The target system must have network connectivity.  
If the target system is *not* connected to the network, and a registration is attempted, the registration fails silently. The Auto Registration process is attempted again during the next reboot.

## Managing the Auto Registration SMF Service

The Auto Registration feature is managed by an SMF service. This service is enabled by default. You can manually disable the service by using the `regadm disable` command. The following procedures describe how to manage the Auto Registration SMF service.

---

**Note** – The Auto Registration SMF service runs in a global zone *only*.

---

TABLE 17-1 Managing the Auto Registration SMF Service (Task Map)

| Task                                                     | Instructions                                                                                                                                                                 | For Information                                                                                |
|----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| Display the status of the Auto Registration SMF service. | You can use the <code>regadm</code> command to determine if the Auto Registration SMF service is running. Note that the service is enabled by default.                       | <a href="#">“How to Determine If the Auto Registration SMF Service Is Enabled” on page 332</a> |
| Disable the Auto Registration SMF service.               | If you do not want to register, or you want to postpone registration, you can optionally disable the Auto Registration SMF service by using the <code>regadm</code> command. | <a href="#">“How to Disable the Auto Registration SMF Service” on page 332</a>                 |
| Enable the Auto Registration SMF service.                | If you previously disabled the Auto Registration SMF service, you must enable the service prior to registering your system.                                                  | <a href="#">“How to Enable the Auto Registration SMF Service” on page 332</a>                  |

## ▼ How to Determine If the Auto Registration SMF Service Is Enabled

- 1 **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

- 2 **Display the current status of the Auto Registration service.**

```
regadm status
Solaris Auto-Registration is currently enabled
```

## ▼ How to Disable the Auto Registration SMF Service

The Auto Registration feature enabled by default. You can manually disable the service by using the `regadm disable` command.

- 1 **Become superuser or assume an equivalent role.**

- 2 **Display the current status of the Auto Registration service.**

```
regadm status
Solaris Auto-Registration is currently enabled
```

- 3 **Disable the Auto Registration service.**

```
regadm disable
```

- 4 **Verify that the service has been disabled.**

```
regadm status
Solaris Auto-Registration is currently disabled
```

## ▼ How to Enable the Auto Registration SMF Service

- 1 **Become superuser or assume an equivalent role.**

- 2 **Display the current status of the Auto Registration service.**

```
regadm status
Solaris Auto-Registration is currently disabled
```

- 3 **Enable the Auto Registration service.**

```
regadm enable
```

**4 Verify that the service has been enabled.**

```
regadm status
Solaris Auto-Registration is currently enabled
```

## Configuring the Auto Registration Feature by Using the regadm Command

The `regadm` command can be used to configure Auto Registration, as well as view configuration details. Note that certain configuration tasks, for example, authentication, must be performed *prior* to registering your system. For more information, see [“Tasks That Are Required Prior to Registration” on page 334](#).

TABLE 17-2 Configuring Auto Registration Properties (Task Map)

| Task                                                                                            | Instructions                                                                                                                                                                                                | For Information                                                                                  |
|-------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Display the current Auto Registration configuration.                                            | Use the <code>regadm list</code> command to display the current Auto Registration configuration. User account credentials, as well as network connectivity information, is displayed by using this command. | <a href="#">“Displaying the Current Auto Registration Configuration” on page 334</a>             |
| Clear the current Auto Registration configuration.                                              | Use the <code>regadm clear</code> command to reset the current Auto Registration configuration. Note that all properties are cleared during this process.                                                   | <a href="#">“How to Clear the Current Auto Registration Configuration” on page 335</a>           |
| Configure HTTP proxy and port settings.                                                         | Use the <code>regadm set</code> command to configure network connectivity information.                                                                                                                      | <a href="#">“How to Configure HTTP Proxy and Port Settings” on page 336</a>                      |
| Configure HTTP proxy authentication settings.                                                   | Use the <code>regadm set</code> command to also configure HTTP proxy authentication settings, per your site's policy and requirements.                                                                      | <a href="#">“How to Configure HTTP Proxy Authentication Settings” on page 337</a>                |
| Authenticate a specific user or account with My Oracle Support by specifying named credentials. | Use the <code>regadm auth</code> command if you want to associate a registered system with a specific user or account.                                                                                      | <a href="#">“How to Authenticate a User by Specifying Named Support Credentials” on page 339</a> |

## Tasks That Are Required Prior to Registration

The process of registering your system by using the `regadm` command includes gathering information and performing certain tasks *prior* to registration. The task flow is as follows:

1. Configure HTTP proxy information, per site policy, if required.
2. Configure HTTP proxy authentication information, per site policy, if required.

For information about how to configure HTTP proxy settings, see [“Configuring Auto Registration Properties” on page 336](#).

3. Select a registration method.

You can register in one of two ways, by specifying named support credentials or anonymously:

- If you want to register by specifying My Oracle Support credentials or by using a Sun Online account, you must provide those support credentials *prior* to registration. For instructions, see [“How to Authenticate a User by Specifying Named Support Credentials” on page 339](#).
- If you are registering anonymously, you can proceed with the registration. See [“How to Register Your System by Using the `regadm` Command” on page 342](#).

## Displaying the Current Auto Registration Configuration

To display the currently authenticated user name and network connectivity information, use the `regadm list` command, as follows:

```
regadm list
```

If you need to reset all of the property settings for a configuration, use the `regadm clear` command. For information, see [“How to Clear the Current Auto Registration Configuration” on page 335](#).

Following are a few examples of various types of Auto Registration configurations.

**EXAMPLE 17-1** Authenticated User (No Proxy)

```
regadm list
My Oracle Support username user@example.com
HTTP Proxy <not configured>
HTTP proxy port <not configured>
HTTP proxy user <not configured>
HTTP proxy password <not configured>
```

**EXAMPLE 17-2** Authenticated User (Open Proxy With No Authentication)

```
regadm list
My Oracle Support username user@example.com
HTTP Proxy webproxy.example.com
HTTP proxy port 8080
HTTP proxy user <not configured>
HTTP proxy password <not configured>
```

**EXAMPLE 17-3** Authenticated User (Proxy That Requires Authentication)

```
regadm list
My Oracle Support username user@example.com
HTTP Proxy webproxy.example.com
HTTP proxy port 8080
HTTP proxy user webuser
HTTP proxy password *****
```

## ▼ How to Clear the Current Auto Registration Configuration

Clearing the current configuration sets all of the Auto Registration properties to their default settings. This operation essentially has the same effect as a factory reset.

- 1 Become superuser or assume an equivalent role.**

- 2 List the current Auto Registration configuration.**

```
regadm list
```

- 3 Clear the current Auto Registration configuration.**

```
regadm clear
```

- 4 Verify that the current configuration was cleared.**

```
regadm list
```

The output now displays the default settings for all property values. The default property values are not configured.

### **Example 17-4** Clearing the Current Auto Registration Configuration

In the following example, the `regadm clear` command is used to clear the property values of the current Auto Registration configuration.

```
regadm list
My Oracle Support username user@example.com
HTTP Proxy webproxy.example.com
HTTP proxy port 8080
HTTP proxy user webuser
```

```

 HTTP proxy password *****
regadm clear
regadm list
 My Oracle Support username <not configured>
 HTTP Proxy <not configured>
 HTTP proxy port <not configured>
 HTTP proxy user <not configured>
 HTTP proxy password <not configured>

```

## Configuring Auto Registration Properties

The following Auto Registration properties can be configured by using the regadm command:

- My Oracle Support or Sun Online account user name and password (Optional)

The user name and password can be configured by using the regadm auth command. The command can be run in either interactive or non-interactive mode. For instructions, see [“How to Authenticate a User by Specifying Named Support Credentials” on page 339](#).

- HTTP proxy name (Optional)
- HTTP proxy port (Optional)

By default, a direct network connection is used during the registration process, unless you specify otherwise *prior* to registering. You can optionally configure HTTP proxy and port settings by using the regadm set command. For instructions, see [“How to Configure HTTP Proxy and Port Settings” on page 336](#).

---

**Note** – To determine the HTTP proxy name and port information to use, check with your site administrator. To determine whether you need to specify HTTP proxy authentication information, check the policy that is implemented at your site.

---

- HTTP proxy user name and password for authentication (Required *only* if HTTP proxy authentication is implemented at your site)

If your site's policy specifies that you must set an HTTP proxy user name and password, you can configure this property prior to registration by using the regadm set command. For instructions, see [“How to Configure HTTP Proxy Authentication Settings” on page 337](#).

Previously configured HTTP proxy and port settings can be changed or reset to a direct network connection by using the regadm clear command. For instructions, see [“How to Clear the Current Auto Registration Configuration” on page 335](#).

### ▼ How to Configure HTTP Proxy and Port Settings

If you previously registered your system by using a direct network connection, or if you are performing a new registration, you might need to configure HTTP proxy and port settings prior to registering your system.



**1 Become superuser or assume an equivalent role.****2 Set the HTTP proxy.**

```
regadm set -n http_proxy -v webproxy-host
```

The web proxy might be an IP address, a non-fully-qualified host name, or a fully-qualified host name.

**3 Set the HTTP port.**

```
regadm set -n http_port -v port-number
```

**▼ How to Configure HTTP Proxy Authentication Settings**

If your site implements HTTP proxy authentication, you must also configure HTTP proxy authentication settings (user name and password). The password must be specified as a single line in a text file and then passed to the regadm command. The following procedure describes how to protect the password file while it is in use. This file should be removed immediately after running the regadm command.

The regadm command accepts a password file in two cases:

- When you are setting a user name and password by using the regadm auth command.
- When you are setting the HTTP proxy user name and password by using the regadm set command.

**1 Become superuser or assume an equivalent role.****2 Set the HTTP proxy user name.**

```
regadm set -n http_proxy_user -v webuser
```

where *webuser* is the site-supplied value to use for authenticating a user with the HTTP proxy server.

**3 Create a temporary password file.**

```
touch pwfile
```

**4 Secure permissions on the file.**

```
chmod 600 pwfile
```

**5 Using a text editor, edit the temporary password file to insert a one-line entry that contains your password information.****6 Set the HTTP password authentication information.**

```
regadm set -n http_proxy_pw -v pwfile
```

where *pwfile* is the password file that you created in Step 3.

## 7 Remove the password file.

```
rm pwfile
```

### Example 17-5 Configuring HTTP Proxy Authentication Settings

The following example shows how to configure the HTTP authentication properties for a site that requires an HTTP user name and password for authentication purposes.

```
regadm set -n http_proxy_user -v webuser
touch myhttppasswd
chmod 600 myhttppasswd
vi myhttppasswd
.
.
.

regadm set -n http_proxy_pw -v myhttppasswd
rm myhttppasswd
```

### Example 17-6 Switching From an HTTP Proxy Network Configuration to a Direct Connection

To switch network connectivity to a direct connection, you must set all of the HTTP proxy settings to a null value (not configured), as shown in the following example:

```
regadm set -n http_proxy ""
regadm set -n http_proxy_port ""
regadm set -n http_proxy_user ""
regadm set -n http_proxy_pw ""
regadm list
 My Oracle Support username <not configured>
 HTTP Proxy <not configured>
 HTTP proxy port <not configured>
 HTTP proxy user <not configured>
 HTTP proxy password <not configured>
```

## Authenticating a User by Specifying Named Support Credentials

Before a system's products can be registered, the user or account must be authenticated by specifying valid My Oracle Support credentials or valid Sun Online account credentials. If no support credentials are provided, the system is registered with an anonymous My Oracle Support account. You can use the regadm command with the auth subcommand, either interactively or non-interactively, to authenticate yourself with the Oracle Product Registration System.

---

**Note** – Authentication must take place *prior* to registering by using the `regadm register` command.

---

The `regadm auth` command is used with the `-u` and `-p` options, as follows:

```
regadm auth -u username -p pwfile
```

The `-u` option specifies a valid user name. The `-p` option specifies a file that contains password information for the user account.

The password file is a single-line file that you create and then add the credentials that are to be passed to the `regadm` command. The contents of the password file are private. For security reasons, remove the file immediately after the information is passed to the `regadm` command and a successful authentication has occurred.

---

**Note** – A successful authentication persists across reboots and only needs to be provided one time. You can authenticate before or during an installation, or after an installation by using the `regadm auth` command. The password is never stored on the system. Only an opaque token is maintained after a successful authentication.

---

## ▼ How to Authenticate a User by Specifying Named Support Credentials

The `regadm auth` command can be used to authenticate a user with My Oracle Support, either interactively or non-interactively. The following procedure describes how to use the command non-interactively. The examples that follow show how to run the command in both modes.

- 1 **Become superuser or assume an equivalent role.**
- 2 **Create a temporary password file:**

```
touch pwfile
```
- 3 **Secure permissions on the file:**

```
chmod 600 pwfile
```
- 4 **Using a text editor, edit the temporary password file that you just created. Insert a one-line entry that contains your password information.**
- 5 **Authenticate with My Oracle Support by specifying your named support credentials:**

```
regadm auth -u user@example.com -p pwfile
```

where `user@example.com` is either the My Oracle Support or Sun Online account user name, and `pwfile` is the password file that you created in Step 2.

A message indicating that the authentication succeeded or failed is displayed.

**6 Proceed as follows:**

- **If the authentication succeeds, you can proceed with the registration.**  
For instructions, see [“How to Register Your System by Using the `regadm` Command” on page 342.](#)
- **If the authentication fails, refer to [“Troubleshooting Auto Registration Issues” on page 345.](#)**

**7 Remove the temporary password file:**

```
rm pwwfile
```

**Example 17-7 Authenticating a User With Named Support Credentials (Non-Interactive)**

```
touch mospwfile
chmod 600 mospwfile
vi mospwfile
.
.
.
regadm auth -u jdoe@mycompany.com -p mospwfile
Authenticating "jdoe@mycompany.com" with My Oracle Support
Authentication successful
```

You can now register this system with My Oracle Support via "regadm register"

```
rm mospwfile
```

**Example 17-8 Authenticating a User With Named Support Credentials (Interactive)**

```
regadm auth
Enter your My Oracle Support username: jdoe@mycompany.com
Enter password:
Enter password again:
Authenticating "jdoe@mycompany.com" with My Oracle Support
Authentication successful
```

You can now register this system with My Oracle Support via "regadm register"

## Registering Your System by Using the `regadm` Command

| Task                                                       | Instructions                                                                                                                                                                                                                                                                                                          | For Information                                                                                         |
|------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------|
| 1. (Optional) Perform a dry run of a product registration. | Use the <code>regadm register</code> command with the <code>-d</code> option to display service tag information about products that are yet to be registered, prior to actually registering. The output of the command displays the data that will be transmitted to My Oracle Support during an actual registration. | <a href="#">“How to Display Registration Information About Products Not Yet Registered” on page 341</a> |
| 2. Perform a product registration.                         | Use the <code>regadm register</code> command to explicitly register your system, independent of an installation or upgrade.                                                                                                                                                                                           | <a href="#">“How to Register Your System by Using the <code>regadm</code> Command” on page 342</a>      |

### ▼ How to Display Registration Information About Products Not Yet Registered

Using the dry run function of the `regadm register` command enables you to display service tag information about products that are not yet registered. This function can be useful for system administrators who want to know what information will be sent to My Oracle Support prior to actually registering. Service tags enable automatic discovery of the products that are installed on your system. When you use the `register -d` command option, the service tag data that is to be transmitted to the Oracle Product Registration System is merely displayed.

---

**Note** – To use the dry run command option you must be on a system that has network connectivity.

---

- 1 **Become superuser or assume an equivalent role.**
- 2 **Type the `regadm register` command with the dry run command option, as follows:**

```
regadm register -d
```

The output of the command is displayed in XML file format.

#### **Example 17–9** Displaying Registration Information About Products Not Yet Registered

The following example shows the output from running the `regadm register` command with the `-d` option. For brevity, the command output has been truncated.

```
regadm register -d
Registering with My Oracle Support using username user@example.com
<?xml version="1.0"?>
<st1:request>
 <agent>
 <agent_urn>urn:st:5eabd994-12d0-6c6c-bf6e-a2cca89a0287</agent_urn>
 <agent_version>1.1.5.</agent_version>
 <registry_version>1.1.5</registry_version>
 <system_info>
 <system>SunOS</system>
 <host>hs-t1000-1</host>
 <release>5.10</release>
 <architecture>sparc</architecture>
 <platform>SUNW,Sun-Fire-T1000::Generic_141444-09</platform>
 .
 .
 </system_info>
 </agent>
</st1:request>

<?xml version="1.0"?>
<st1:request>
 <service_tag>
 <instance_urn>urn:st:fd96de223-e190-6501-ff8d-dd76f1b708f5</instance_urn>
 <product_name>Solaris 10 Operating System</product_name>
 .
 .
 .
</st1:request>
```

## ▼ How to Register Your System by Using the regadm Command

The `regadm register` command is used to register each product in the Service Tag Registry with My Oracle Support. Registration is accomplished by using the current authentication credentials. If no credentials are provided, the registration is performed anonymously by using an anonymous My Oracle Support account.

**Before You Begin** Be sure to check [“Tasks That Are Required Prior to Registration”](#) on page 334 for prerequisite instructions.

Consider also the following information:

- If you previously registered your system anonymously, and you now want to register by specifying named support credentials, you must authenticate those credentials by using the `regadm auth` command. The process of authentication is a separate task that must be performed *prior* to registration. For instructions, see [“How to Authenticate a User by Specifying Named Support Credentials” on page 339](#).
- If you previously registered your system by specifying named support credentials, and you now want to register anonymously, you must clear the current Auto Registration configuration *prior* to registration. For instructions, see [“How to Clear the Current Auto Registration Configuration” on page 335](#).

---

**Note** – The `clear` subcommand clears all configured properties, including any previously configured network connectivity information.

---

- If you previously registered your system by using a direct network connection, and you now want to specify HTTP proxy and port information, you must configure those properties by using the `regadm set` command *prior* to registration. For instructions, see [“How to Configure HTTP Proxy and Port Settings” on page 336](#).

### 1 Become superuser or assume an equivalent role.

### 2 List the current Auto Registration configuration.

```
regadm list
```

If the output of the `regadm list` command does not display any configured Auto Registration properties (user name), the registration will be anonymous. Otherwise, the registration is associated with the user name that is displayed. If you are registering your system anonymously, but you prefer to provide support credentials, see [“How to Authenticate a User by Specifying Named Support Credentials” on page 339](#) before proceeding.

If the output of the `regadm list` command does not display any HTTP proxy or port information, the registration will be performed by using a direct network connection. Otherwise, the existing HTTP property settings are used. If you need to reconfigure Auto Registration properties for network connectivity, see [“How to Configure HTTP Proxy and Port Settings” on page 336](#) before proceeding.

### 3 Register your system.

```
regadm register
```

---

**Note** – The output of the `regadm register` command varies, depending on how the Auto Registration feature is configured on your system.

---

**Example 17-10** Registering Your System Anonymously

The following example shows an anonymous registration on a system where no named support credentials were configured prior to registration. The registration also uses a direct network connection.

```
regadm list
 My Oracle Support username <not configured>
 HTTP Proxy <not configured>
 HTTP proxy port <not configured>
 HTTP proxy user <not configured>
 HTTP proxy password <not configured>
regadm register
Registered Solaris 10 Operating System
Registration Successful
```

**Example 17-11** Registering Your System by Specifying Named Support Credentials

The following example shows a registration on a system where support credentials and HTTP proxy and port information were configured prior to registration.

```
regadm list
 My Oracle Support username user@example.com
 HTTP Proxy webproxy.example.com
 HTTP proxy port 8080
 HTTP proxy user <not configured>
 HTTP proxy password <not configured>
regadm register
Registering with My Oracle Support using username user@example.com
Registered Solaris 10 Operating System
Registration successful
```

## Obtaining Information About Currently Registered Products

To obtain service tag information for products that are currently registered with the Auto Registration feature, go to <https://inventory.sun.com/inventory/data.jsp>.

To access service tag information for your registered products, you must first log in with either your My Oracle Support account credentials or your Sun Online account credentials.

---

**Note** – Be aware that the inventory web site provides information about discovery and registration methods that are not associated with the new Auto Registration feature. These other registration methods are used on a voluntary basis and registration is not mandatory. Starting from Oracle Solaris 10 9/10, Auto Registration occurs upon the first reboot after an installation or upgrade.

---



For general information about how your system's assets are managed with service tags, go to <http://www.oracle.com/technetwork/server-storage/solaris/oracle-service-tag-faq-418684.html>.

## Troubleshooting Auto Registration Issues

The more common errors that you might encounter while attempting to authenticate or register your system can be resolved by following basic troubleshooting strategies. The following are issues that you might need to troubleshoot when using the `regadm` command:

### Authentication Failed

If you attempt to authenticate yourself with My Oracle Support by using the `regadm auth` command, and the authentication fails, the following message is displayed after you type your user name and password:

```
Authentication failed
```

To troubleshoot this issue, verify that you are using the correct account credentials (My Oracle Support user name and password, or Sun Online account user name and password). To verify your account settings, go to <https://reg.sun.com/accounthelp>.

For answers to common questions and information about how to gain initial access to My Oracle Support, go to <http://www.oracle.com/us/support/044753.html#SS6>.

### Cannot Send HTTP Proxy Password

For security reasons, the HTTP proxy password must be located in a file rather than specified on the command line. If the password has not been specified in a file, the `regadm` command cannot locate the password. Consequently, the following message is displayed when you attempt to set your password:

```
regadm set -n http_proxy_pw -v secret1
Cannot find file 'secret1'
```

For instructions on setting the HTTP proxy password, see “[How to Configure HTTP Proxy Authentication Settings](#)” on page 337.

### HTTP Communications Error

If a problem exists with your HTTP configuration, the following error message might be displayed:

HTTP Communications error with Oracle Support Center

Please check your HTTP settings and retry your request.

There can be one or more reasons that this error message is displayed. To troubleshoot the problem, follow these steps:

1. Verify that your HTTP configuration is correct by typing the following command:

```
regadm list
```

2. If your configuration is correct, check the following information:
  - Does your system have access to the public Internet, either directly or through an HTTP proxy? If not, you cannot use the `regadm` command to administer the Auto Registration feature.
  - Is an HTTP proxy required at your site? If so, ensure that the HTTP proxy host and port have been configured correctly. For instructions, see [“How to Configure HTTP Proxy and Port Settings” on page 336](#).
  - Is authenticated proxy access required at your site? If so, verify that the proxy user name and password have been configured and typed correctly. For instructions, see [“How to Configure HTTP Proxy Authentication Settings” on page 337](#).
3. Ensure that your HTTP proxy host setting can be resolved to a valid web proxy host that is up and running and has a name that is resolvable.

For example, if your configured HTTP proxy host is `webproxy.example.com`, it should be both resolvable and able to respond to a request from the `ping` utility, as follows:

```
getent hosts webproxy.example.com
192.168.1.1 webproxy.example.com
ping webproxy.example.com
webproxy.example.com is alive
```

For more information, see the [`getent\(1M\)`](#) and [`ping\(1M\)`](#) man pages.

4. If the previous troubleshooting steps do not work, verify your settings with your site administrator.

# Managing Services (Overview)

---

This chapter provides an overview of the Service Management Facility (SMF). In addition, information that is related to run levels is provided.

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

- “Introduction to SMF” on page 347
- “SMF Concepts” on page 349
- “SMF Administrative and Programming Interfaces” on page 353
- “SMF Components” on page 354
- “SMF Compatibility” on page 356
- “Run Levels” on page 356
- “/etc/inittab File” on page 358

For information on the procedures associated with SMF, see “Managing Services (Task Map)” on page 361. For information on the procedures associated with run levels, see “Using Run Control Scripts (Task Map)” on page 376.

## Introduction to SMF

SMF provides an infrastructure that augments the traditional UNIX start-up scripts, `init` run levels, and configuration files. SMF provides the following functions:

- Automatically restarts failed services in dependency order, whether they failed as the result of administrator error, software bug, or were affected by an uncorrectable hardware error. The dependency order is defined by dependency statements.
- Makes services objects that can be viewed, with the new `svcs` command, and managed, with `svcadm` and `svccfg` commands. You can also view the relationships between services and processes using `svcs -p`, for both SMF services and legacy `init.d` scripts.
- Makes it easy to backup, restore, and undo changes to services by taking automatic snapshots of service configurations.

- Makes it easy to debug and ask questions about services by providing an explanation of why a service isn't running by using `svcs -x`. Also, this process is eased by individual and persistent log files for each service.
- Allows for services to be enabled and disabled using `svcadm`. These changes can persist through upgrades and reboots. If the `-t` option is used, the changes are temporary.
- Enhances the ability of administrators to securely delegate tasks to non-root users, including the ability to modify properties and enable, disable, or restart services on the system.
- Boots faster on large systems by starting services in parallel according to the dependencies of the services. The opposite process occurs during shutdown.
- Allows you to customize the boot console output to either be as quiet as possible, which is the default, or to be verbose by using `boot -m verbose`.
- Preserves compatibility with existing administrative practices wherever possible. For example, most customer and ISV-supplied rc scripts still work as usual.

*Dependency statements* define the relationships between services. These relationships can be used to provide precise fault containment by restarting only those services that are directly affected by a fault, rather than restarting all of the services. Another advantage of dependency statements is that the statements allow for scalable and reproducible initialization processes. In addition, by defining all of the dependencies, you can take advantage of modern, highly parallel machines, because all independent services can be started in parallel.

SMF defines a set of actions that can be invoked on a service by an administrator. These actions include enable, disable, refresh, restart, and maintain. Each service is managed by a service restarter which carries out the administrative actions. In general, the restarters carry out actions by executing methods for a service. Methods for each service are defined in the service configuration repository. These methods allow the restarter to move the service from one state to another state.

The service configuration repository provides a per-service snapshot at the time that each service is successfully started so that fallback is possible. In addition, the repository provides a consistent and persistent way to enable or disable a service, as well as a consistent view of service state. This capability helps you debug service configuration problems.

## Changes in Behavior When Using SMF

Most of the features that are provided by SMF happen behind the scenes, so users are not aware of them. Other features are accessed by new commands. Here is a list of the behavior changes that are most visible.

- The boot process creates many fewer messages now. Services do not display a message by default when they are started. All of the information that was provided by the boot messages can now be found in a log file for each service that is in `/var/svc/LOG`. You can use the `svcs`

command to help diagnose boot problems. In addition, you can use the `-v` option to the boot command, which generates a message when each service is started during the boot process.

- Since services are automatically restarted if possible, it may seem that a process refuses to die. If the service is defective, the service will be placed in maintenance mode, but normally a service is restarted if the process for the service is killed. The `svcadm` command should be used to stop the processes of any SMF service that should not be running.
- Many of the scripts in `/etc/init.d` and `/etc/rc*.d` have been removed. The scripts are no longer needed to enable or disable a service. Entries from `/etc/inittab` have also been removed, so that the services can be administered using SMF. Scripts and `inittab` entries that are provided by an ISV or are locally developed will continue to run. The services may not start at exactly the same point in the boot process, but they are not started before the SMF services, so that any service dependencies should be OK.

## SMF Concepts

This section presents terms and their definitions within the SMF framework. These terms are used throughout the documentation. To grasp SMF concepts, an understanding of these terms is essential.

### SMF Service

The fundamental unit of administration in the SMF framework is the *service instance*. Each SMF service has the potential to have multiple versions of it configured. As well, multiple instances of the same version can run on a single Oracle Solaris system. An *instance* is a specific configuration of a service. A web server is a service. A specific web server daemon that is configured to listen on port 80 is an instance. Each instance of the web server service could have different configuration requirements. The service has system-wide configuration requirements, but each instance can override specific requirements, as needed. Multiple instances of a single service are managed as child objects of the service object.

Services are not just the representation for standard long-running system services such as `in.dhcpd` or `in.nfsd`. Services also represent varied system entities that include ISV applications such as Oracle software. In addition, a service can include less traditional entities such as the following:

- A physical network device
- A configured IP address
- Kernel configuration information
- Milestones that correspond to system init state, such as the multiuser run level

Generically, a service is an entity that provides a list of capabilities to applications and other services, local and remote. A service is dependent on an implicitly declared list of local services.

A *milestone* is a special type of service. Milestone services represent high-level attributes of the system. For example, the services which constitute run levels S, 2, and 3 are each represented by milestone services.

## Service Identifiers

Each service instance is named with a Fault Management Resource Identifier or FMRI. The FMRI includes the service name and the instance name. For example, the FMRI for the `rlogin` service is `svc:/network/login:rlogin`, where `network/login` identifies the service and `rlogin` identifies the service instance.

Equivalent formats for an FMRI are as follows:

- `svc://localhost/system/system-log:default`
- `svc:/system/system-log:default`
- `system/system-log:default`

In addition, some SMF commands can use the following FMRI format:

`svc:/system/system-log`. Some commands infer what instance to use, when there is no ambiguity. See the SMF command man pages, such as [svcadm\(1M\)](#) or [svcs\(1\)](#), for instructions about which FMRI formats are appropriate.

The service names usually include a general functional category. The categories include the following:

- `application`
- `device`
- `milestone`
- `network`
- `platform`
- `site`
- `system`

Legacy `init.d` scripts are also represented with FMRI that start with `lrc` instead of `svc`, for example: `lrc:/etc/rcS_d/S35cacheos_sh`. The legacy services can be monitored using SMF. However, you cannot administer these services.

When booting a system for the first time with SMF, services listed in `/etc/inetd.conf` are automatically converted into SMF services. The FMRI for these services are slightly different. The syntax for a converted `inetd` services is:

```
network/<service-name>/<protocol>
```

In addition, the syntax for a converted service that uses the RPC protocol is:

```
network/rpc-<service-name>/rpc_<protocol>
```

Where *<service-name>* is the name defined in `/etc/inetd.conf` and *<protocol>* is the protocol for the service. For instance, the FMRI for the `rpc.cmsd` service is `network/rpc-100068_2-5/rpc_udp`.

## Service States

The `svcs` command displays the state, start time, and FMRI of service instances. The state of each service is one of the following:

- `degraded` – The service instance is enabled, but is running at a limited capacity.
- `disabled` – The service instance is not enabled and is not running.
- `legacy_run` – The legacy service is not managed by SMF, but the service can be observed. This state is only used by legacy services.
- `maintenance` – The service instance has encountered an error that must be resolved by the administrator.
- `offline` – The service instance is enabled, but the service is not yet running or available to run.
- `online` – The service instance is enabled and has successfully started.
- `uninitialized` – This state is the initial state for all services before their configuration has been read.

## SMF Manifests

An SMF *manifest* is an XML file that contains a complete set of properties that are associated with a service or a service instance. The files are stored in `/var/svc/manifest`. Manifests should not be used to modify the properties of a service. The service configuration repository is the authoritative source of configuration information. To incorporate information from the manifest into the repository, you must either run `svccfg import` or allow the service to import the information during a system boot.

See the [service\\_bundle\(4\)](#) man page for a complete description of the contents of the SMF manifests. If you need to change the properties of a service, see the [svccfg\(1M\)](#) or [inetadm\(1M\)](#) man pages.

## SMF Profiles

An SMF *profile* is an XML file that lists a set of service instances and whether each should be enabled or disabled. Some profiles which are delivered with the Oracle Solaris release include:

- `/var/svc/profile/generic_open.xml` – This profile enables the standard services that have been started by default in earlier Solaris releases.
- `/var/svc/profile/generic_limited_net.xml` – This profile disables many of the internet services that have been started by default in earlier Solaris releases. The `network/ssh` service is enabled to provide network connectivity.
- `/var/svc/profile/ns_*.xml` – These profiles enable services associated with the name service that is configured to run on the system.
- `/var/svc/profile/platform_*.xml` – These profiles enable services associated with particular hardware platforms.

During the first boot after a new installation or an upgrade to the Oracle Solaris OS, some Solaris profiles are automatically applied. To be specific, the `/var/svc/profile/generic.xml` profile is applied. This file is usually symbolically linked to `generic_open.xml` or `generic_limited_net.xml`. Also, if a profile called `site.xml` is in `/var/svc/profile` during the first boot or is added between boots, the contents of this profile are applied. By using the `site.xml` profile, the initial set of enabled services may be customized by the administrator.

For more information about using profiles, see [“How to Apply an SMF Profile” on page 370](#).

## Service Configuration Repository

The *service configuration repository* stores persistent configuration information as well as SMF runtime data for services. The repository is distributed among local memory and local files. SMF is designed so that eventually, service data can be represented in the network directory service. The network directory service is not yet available. The data in the service configuration repository allows for the sharing of configuration information and administrative simplicity across many Solaris instances. The service configuration repository can only be manipulated or queried using SMF interfaces. For more information about manipulating and accessing the repository, see the `svccfg(1M)` and `svcprop(1)` man pages. The service configuration repository daemon is covered in the `svc.configd(1M)` man page. The service configuration library is documented in the `libscf(3LIB)` man page.



## SMF Repository Backups

SMF automatically takes the following backups of the repository:

- The boot backup is taken immediately before the first change to the repository is made during each system startup.
- The `manifest_import` backup occurs after `svc:/system/manifest-import:default` completes, if it imported any new manifests or ran any upgrade scripts.

Four backups of each type are maintained by the system. The system deletes the oldest backup, when necessary. The backups are stored as `/etc/svc/repository-type-YYYYMMDD_HHMMSS`, where `YYYYMMDD` (year, month, day) and `HHMMSS` (hour, minute, second), are the date and time when the backup was taken. Note that the hour format is based on a 24-hour clock.

You can restore the repository from these backups, if an error occurs. To do so, use the `/lib/svc/bin/restore_repository` command. For more information, see [“How to Repair a Corrupt Repository” on page 379](#).

## SMF Snapshots

The data in the service configuration repository includes *snapshots*, as well as a configuration that can be edited. Data about each service instance is stored in the snapshots. The standard snapshots are as follows:

- `initial` – Taken on the first import of the manifest
- `running` – Used when the service methods are executed
- `start` – Taken at the last successful start

The SMF service always executes with the `running` snapshot. This snapshot is automatically created if it does not exist.

The `svcadm refresh` command, sometimes followed by the `svcadm restart` command, makes a snapshot active. The `svccfg` command is used to view or revert to instance configurations in a previous snapshot. See [“How to Revert to Another SMF Snapshot” on page 368](#) for more information.

## SMF Administrative and Programming Interfaces

This section introduces the interfaces that are available when you use SMF.

## SMF Command-Line Administrative Utilities

SMF provides a set of command-line utilities that interact with SMF and accomplish standard administrative tasks. The following utilities can be used to administer SMF.

TABLE 18-1 Service Management Facility Utilities

| Command Name          | Function                                                                                                                      |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <code>inetadm</code>  | Provides the ability to observe or configure services controlled by <code>inetd</code>                                        |
| <code>svcadm</code>   | Provides the ability to perform common service management tasks, such as enabling, disabling, or restarting service instances |
| <code>svccfg</code>   | Provides the ability to display and manipulate the contents of the service configuration repository                           |
| <code>svccprop</code> | Retrieves property values from the service configuration repository with a output format appropriate for use in shell scripts |
| <code>svcs</code>     | Gives detailed views of the service state of all service instances in the service configuration repository                    |

## Service Management Configuration Library Interfaces

SMF provides a set of programming interfaces that are used to interact with the service configuration repository through the `svc.configd` daemon. This daemon is the arbiter of all requests to the local repository datastores. A set of fundamental interfaces is defined as the lowest level of interaction possible with services in the service configuration repository. The interfaces provide access to all service configuration repository features such as transactions and snapshots.

Many developers only need a set of common tasks to interact with SMF. These tasks are implemented as convenience functions on top of the fundamental services to ease the implementation burden.

## SMF Components

SMF includes a master restarter daemon and delegated restarters.

### SMF Master Restarter Daemon

The `svc.startd` daemon is the master process starter and restarter for the Solaris OS. The daemon is responsible for managing service dependencies for the entire system. The daemon takes on the previous responsibility that `init` held of starting the appropriate `/etc/rc*.d`

scripts at the appropriate run levels. First, `svc.startd` retrieves the information in the service configuration repository. Next, the daemon starts services when their dependencies are met. The daemon is also responsible for restarting services that have failed and for shutting down services whose dependencies are no longer satisfied. The daemon keeps track of service state through an operating system view of availability through events such as process death.

## SMF Delegated Restarters

Some services have a set of common behaviors on startup. To provide commonality among these services, a delegated restarter might take responsibility for these services. In addition, a delegated restarter can be used to provide more complex or application-specific restarting behavior. The delegated restarter can support a different set of methods, but exports the same service states as the master restarter. The restarter's name is stored with the service. A current example of a delegated restarter is `inetd`, which can start Internet services on demand, rather than having the services always running.

## SMF and Booting

SMF provides new methods for booting a system. For instance:

- There is a additional system state which is associated with the `all` milestone. With the `all` milestone, all of the services with a defined dependency on the `multi-user-server` milestone are started, as well as any services that do not have a defined dependency. If you have added services, such as third party products, they may not be started automatically unless you use the following command:  

```
ok boot -m milestone=all
```
- When booting a system, you can choose to use the verbose option to see more messages. By default, the system will not display these messages. To boot in the verbose mode, use the following command:  

```
ok boot -mverbose
```
- There is a new system state which is associated with the `none` milestone. Only `init`, `svc.startd` and `svc.configd` are started if you boot a system using this milestone. This state can be very useful for debugging booting problems. In particular, debugging any problems with the configuration of SMF services is made simpler, because none of the services are started. See [“How to Boot Without Starting Any Services” on page 382](#) for instructions on how to use the `none` milestone.

## SMF Compatibility

While many standard Solaris services are now managed by SMF, the scripts placed in `/etc/rc*.d` continue to be executed on run-level transitions. Most of the `/etc/rc*.d` scripts that were included in previous Solaris releases have been removed as part of SMF. The ability to continue to run the remaining scripts allows for third-party applications to be added without having to convert the services to use SMF.

In addition, `/etc/inittab` and `/etc/inetd.conf` must be available for packages to amend with postinstall scripts. These are called legacy-run services. The `inetconv` command is run to add these legacy-run services to the service configuration repository. The status of these services can be viewed, but no other changes are supported through SMF. Applications that use this feature will not benefit from the precise fault containment provided by SMF.

Applications converted to utilize SMF should no longer make modifications to the `/etc/inittab` and `/etc/inetd.conf` files. The converted applications will not use the `/etc/rc*.d` scripts. Also, the new version of `inetd` does not look for entries in `/etc/inetd.conf`.

## Run Levels

A system's *run level* (also known as an *init state*) defines what services and resources are available to users. A system can be in only one run level at a time.

The Solaris OS has eight run levels, which are described in the following table. The default run level is specified in the `/etc/inittab` file as run level 3.

TABLE 18–2 Solaris Run Levels

| Run Level | Init State                                | Type        | Purpose                                                                                                                                     |
|-----------|-------------------------------------------|-------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| 0         | Power-down state                          | Power-down  | To shut down the operating system so that it is safe to turn off power to the system.                                                       |
| s or S    | Single-user state                         | Single-user | To run as a single user with some file systems mounted and accessible.                                                                      |
| 1         | Administrative state                      | Single-user | To access all available file systems. User logins are disabled.                                                                             |
| 2         | Multiuser state                           | Multiuser   | For normal operations. Multiple users can access the system and all file system. All daemons are running except for the NFS server daemons. |
| 3         | Multiuser level with NFS resources shared | Multiuser   | For normal operations with NFS resources shared. This is the default run level for the Solaris OS.                                          |

TABLE 18–2 Solaris Run Levels (Continued)

| Run Level | Init State                  | Type       | Purpose                                                                                                                                                                    |
|-----------|-----------------------------|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 4         | Alternative multiuser state |            | Not configured by default, but available for customer use.                                                                                                                 |
| 5         | Power-down state            | Power-down | To shut down the operating system so that it is safe to turn off power to the system. If possible, automatically turns off power on systems that support this feature.     |
| 6         | Reboot state                | Reboot     | To shut down the system to run level 0, and then reboot to multiuser level with NFS resources shared (or whatever level is the default in the <code>init</code> tab file). |

In addition, the `svcadm` command can be used to change the run level of a system, by selecting a milestone at which to run. The following table shows which run level corresponds to each milestone.

TABLE 18–3 Solaris Run Levels and SMF Milestones

| Run Level | SMF Milestone FMRI                  |
|-----------|-------------------------------------|
| S         | milestone/single-user:default       |
| 2         | milestone/multi-user:default        |
| 3         | milestone/multi-user-server:default |

## When to Use Run Levels or Milestones

Under most circumstances, using the `init` command with a run level to change the system state is sufficient. Using milestones to change system state can be confusing and can lead to unexpected behavior. In addition, the `init` command allows for the system to be shutdown, so `init` is the best command for changing system state.

However, booting a system using the `none` milestone, can be very useful when debugging startup problems. There is no equivalent run level to the `none` milestone. See [“How to Boot Without Starting Any Services” on page 382](#) for specific instructions.

## Determining a System's Run Level

Display run level information by using the `who -r` command.

```
$ who -r
```

Use the `who -r` command to determine a system's current run level for any level.

**EXAMPLE 18-1** Determining a System's Run Level

This example displays information about a system's current run level and previous run levels.

```
$ who -r
. run-level 3 Dec 13 10:10 3 0 S
$
```

| Output of <code>who -r</code> command | Description                                                                                |
|---------------------------------------|--------------------------------------------------------------------------------------------|
| run-level 3                           | Identifies the current run level                                                           |
| Dec 13 10:10                          | Identifies the date of last run level change                                               |
| 3                                     | Also identifies the current run level                                                      |
| 0                                     | Identifies the number of times the system has been at this run level since the last reboot |
| S                                     | Identifies the previous run level                                                          |

## /etc/inittab File

When you boot the system or change run levels with the `init` or `shutdown` command, the `init` daemon starts processes by reading information from the `/etc/inittab` file. This file defines these important items for the `init` process:

- That the `init` process will restart
- What processes to start, monitor, and restart if they terminate
- What actions to take when the system enters a new run level

Each entry in the `/etc/inittab` file has the following fields:

*id*:*rstate*:*action*:*process*

The following table describes the fields in an `inittab` entry.

**TABLE 18-4** Fields Descriptions for the `inittab` File

| Field          | Description                                                                                                                                                                                                                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>id</i>      | Is a unique identifier for the entry.                                                                                                                                                                                                                                                                          |
| <i>rstate</i>  | Lists the run levels to which this entry applies.                                                                                                                                                                                                                                                              |
| <i>action</i>  | Identifies how the process that is specified in the process field is to be run. Possible values include: <code>sysinit</code> , <code>boot</code> , <code>bootwait</code> , <code>wait</code> , and <code>respawn</code> .<br>For a description of the other action keywords, see <a href="#">inittab(4)</a> . |
| <i>process</i> | Defines the command or script to execute.                                                                                                                                                                                                                                                                      |

**EXAMPLE 18-2** Default inittab File

The following example shows a default `inittab` file that is installed with the Solaris release. A description for each line of output in this example follows.

```
ap::sysinit:/sbin/autopush -f /etc/iu.ap (1)
sp::sysinit:/sbin/soconfig -f /etc/sock2path (2)
smf::sysinit:/lib/svc/bin/svc.startd >/dev/msglog 2<>/dev/msglog (3)
p3:s1234:powerfail:/usr/sbin/shutdown -y -i5 -g0 >/dev/msglog 2<>/dev/... (4)
```

1. Initializes STREAMS modules
2. Configures socket transport providers
3. Initializes the master restarter for SMF
4. Describes a power fail shutdown

## What Happens When the System Is Brought to Run Level 3

1. The `init` process is started and reads the `/etc/default/init` file to set any environment variables. By default, only the `TIMEZONE` variable is set.
2. Then, `init` reads the `inittab` file and does the following:
  - a. Executes any process entries that have `sysinit` in the action field so that any special initializations can take place before users login.
  - b. Passes the startup activities to `svc.startd`.

For a detailed description of how the `init` process uses the `inittab` file, see [init\(1M\)](#).





# Managing Services (Tasks)

---

This chapter covers the tasks required to manage and monitor the Service Management Facility (SMF). In addition, information that is related to managing run level scripts is provided. The following topics are covered:

- “Managing Services (Task Map)” on page 361
- “Monitoring SMF Services” on page 362
- “Managing SMF Services” on page 365
- “Configuring SMF Services” on page 371
- “Using Run Control Scripts” on page 376
- “Troubleshooting the Service Management Facility” on page 379

## Managing Services (Task Map)

The following task map describes the procedures that are needed to use SMF.

| Task                                      | Description                                                                                                                                              | For Instructions                                                                             |
|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Display the status of a service instance. | Displays the status of all running service instances.                                                                                                    | <a href="#">“How to List the Status of a Service” on page 362</a>                            |
| Display the service dependents.           | Display the services that are dependent on the specified service.                                                                                        | <a href="#">“How to Show Which Services Are Dependent on a Service Instance” on page 364</a> |
| Display the dependencies of a service.    | Display the services that a specified service is dependent on. This information can be used to help identify what is preventing a service from starting. | <a href="#">“How to Show Which Services a Service Is Dependent On” on page 364</a>           |
| Disable a service instance.               | Turns off a service that is not functioning properly or needs to be off to increase security.                                                            | <a href="#">“How to Disable a Service Instance” on page 366</a>                              |

| Task                                               | Description                                                                                                                                                                                                                                       | For Instructions                                                                                                                                                                                                                                                                         |
|----------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable a service instance                          | Starts a service.                                                                                                                                                                                                                                 | <a href="#">“How to Enable a Service Instance” on page 366</a>                                                                                                                                                                                                                           |
| Restart a service instance.                        | Restart a service, without having to use separate commands to disable and then enable the service.                                                                                                                                                | <a href="#">“How to Restart a Service” on page 367</a>                                                                                                                                                                                                                                   |
| Modify a service instance.                         | Modifies the configuration parameters of a specified service instance.<br><br>Changes a configuration property of a service controlled by <code>inetd</code> .<br><br>Changes the startup options of a service controlled by <code>inetd</code> . | <a href="#">“How to Modify a Service” on page 371</a><br><br><a href="#">“How to Change a Property for an <code>inetd</code> Controlled Service” on page 373</a><br><br><a href="#">“How to Modify a Command-Line Argument for an <code>inetd</code> Controlled Service” on page 374</a> |
| Convert <code>inetd.conf</code> entries.           | Converts <code>inetd</code> services into legacy-run services that can be monitored using SMF.                                                                                                                                                    | <a href="#">“How to Convert <code>inetd.conf</code> Entries” on page 375</a>                                                                                                                                                                                                             |
| Repair a corrupt service configuration repository. | Replaces a corrupt repository with a default version.                                                                                                                                                                                             | <a href="#">“How to Repair a Corrupt Repository” on page 379</a>                                                                                                                                                                                                                         |
| Boot a system without starting any services.       | Boots a system without starting any services so that configuration problems that prevent booting can be fixed.                                                                                                                                    | <a href="#">“How to Boot Without Starting Any Services” on page 382</a>                                                                                                                                                                                                                  |

## Monitoring SMF Services

The following tasks show how to monitor SMF services.

### ▼ How to List the Status of a Service

This procedure can be used to show what services are running.

- **Run the `svcs` command.**

Running this command without any options displays a status report of the service specified by the FMRI.

```
% svcs -l FMRI
```

#### Example 19-1 Showing the Status of the `rlogin` Service

This example shows the status of a service that includes many contracts.

```
% svcs -l network/login:rlogin
fmri svc:/network/login:rlogin
name remote login
enabled true
state online
next_state none
state_time Thu Apr 28 14:10:48 2011
restarter svc:/network/inetd:/default
contract_id 42325 41441 40776 40348 40282 40197 39025 38381 38053\
33697 28625 24652 23689 15352 9889 7194 6576 6360 5387 1475 3015\
6545 6612 9302 9662 10484 16254 19850 22512 23394 25876 26113 27326\
34284 37939 38405 38972 39200 40503 40579 41129 41194
```

### Example 19-2 Showing the Status of the sendmail Service

This example shows the status of a service that includes dependencies.

```
% svcs -l network/smtp:sendmail
fmri svc:/network/smtp:sendmail
name sendmail SMTP mail transfer agent
enabled true
state online
next_state none
state_time Thu Apr 28 14:10:37 2011
restarter svc:/system/svc/restarter:default
contract_id 29462
dependency require_all/refresh file://localhost/etc/nsswitch.conf (-)
dependency require_all/refresh file://localhost/etc/mail/sendmail.cf (-)
dependency optional_all/none svc:/system/system-log (online)
dependency require_all/refresh svc:/system/identity:domain (online)
dependency require_all/refresh svc:/milestone/name-services (online)
dependency require_all/none svc:/network/service (online)
dependency require_all/none svc:/system/filesystem/local (online)
```

### Example 19-3 Showing the Status of all Services

The following command lists all services that are installed on the system as well as the status of each service. The command displays those services that are disabled as well as those that are enabled.

```
% svcs -a
```

### Example 19-4 Showing the Status of Services Controlled by inetd

The following command lists services that are controlled by `inetd`. Each service's FMRI is listed, along with the run state and whether the service is enabled or disabled.

```
% inetadm
```

## ▼ How to Show Which Services Are Dependent on a Service Instance

This procedure shows how to determine which service instances depend on the specified service.

- Display the service dependents.

```
% svcs -D FMRI
```

### Example 19–5 Displaying the Service Instances That Are Dependent on the Multiuser Milestone

The following example shows how to determine which service instances are dependent on the multiuser milestone.

```
% svcs -D milestone/multi-user
STATE STIME FMRI
online Apr_08 svc:/milestone/multi-user-server:default
```

## ▼ How to Show Which Services a Service Is Dependent On

This procedure shows how to determine which services a specified service instance is dependent on.

- Display the service dependencies.

```
% svcs -d FMRI
```

### Example 19–6 Displaying the Service Instances That the Multiuser Milestone Is Dependent On

The following example shows the services instances that the multiuser milestone is dependent on.

```
% svcs -d milestone/multi-user:default
STATE STIME FMRI
disabled Aug_24 svc:/platform/sun4u/sf880drd:default
online Aug_24 svc:/milestone/single-user:default
online Aug_24 svc:/system/utmp:default
online Aug_24 svc:/system/system-log:default
online Aug_24 svc:/system/system-log:default
online Aug_24 svc:/system/rmtmpfiles:default
online Aug_24 svc:/network/rpc/bind:default
online Aug_24 svc:/milestone/name-services:default
online Aug_24 svc:/system/filesystem/local:default
online Aug_24 svc:/system/mdmonitor:default
```

## Managing SMF Services (Task Map)

| Task                                                                                     | Description                                                                                                                                                                                     | For Instructions                                                                                      |
|------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Disable a service instance.                                                              | Stops a running service and prevents the service from restarting.                                                                                                                               | <a href="#">“How to Disable a Service Instance” on page 366</a>                                       |
| Enable a service instance.                                                               | Starts a service. In addition, the service will be restarted during subsequent reboots.                                                                                                         | <a href="#">“How to Enable a Service Instance” on page 366</a>                                        |
| Restarting a service.                                                                    | Stops and starts a service with one command.                                                                                                                                                    | <a href="#">“How to Restart a Service” on page 367</a>                                                |
| Restoring a service in maintenance state.                                                | Shows how to clean up and restart a service that is in maintenance state.                                                                                                                       | <a href="#">“How to Restore a Service That Is in the Maintenance State” on page 367</a>               |
| Revert to a snapshot.                                                                    | Uses a previous snapshot to correct problems with a service.                                                                                                                                    | <a href="#">“How to Revert to Another SMF Snapshot” on page 368</a>                                   |
| Create an profile.                                                                       | Create a profile to disable or enable services as needed.                                                                                                                                       | <a href="#">“How to Create an SMF Profile” on page 369</a>                                            |
| Apply a profile.                                                                         | Uses the information in a profile to disable or enable services as needed.                                                                                                                      | <a href="#">“How to Apply an SMF Profile” on page 370</a>                                             |
| Change the services and their configuration using the <code>net services</code> command. | Uses the information in the <code>generic_limited.xml</code> or <code>generic_open.xml</code> profiles to disable or enable services and make configuration changes to those services, as well. | <a href="#">“Changing Services Offered to the Network with <code>generic*.xml</code>” on page 371</a> |

## Managing SMF Services

This section includes information on managing SMF services.

### Using RBAC Rights Profiles With SMF

You can use RBAC rights profiles to allow users to manage some of the SMF services, without having to give the user root access. The rights profiles define what commands the user can run. For SMF, the following profiles have been created:

- **Service Management:** User can add, delete or modify services.
- **Service Operator:** User can request state changes of any service instance, such as restart and refresh.

For specific information about the authorizations, see the `smf_security(5)` man page. For instructions to assign a rights profile, see [“How to Change the RBAC Properties of a User” in \*System Administration Guide: Security Services\*](#).

## ▼ How to Disable a Service Instance

Use the following procedure to disable a service. The service status change is recorded in the service configuration repository. Once the service is disabled, the disabled state will persist across reboots. The only way to get the service running again is to enable it.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 Check the dependents of the service you want to disable.

If this service has dependents that you need, then you cannot disable this service.

```
svcs -D FMRI
```

### 3 Disable the service.

```
svcadm disable FMRI
```

#### Example 19-7 Disabling the rlogin Service

The output from the first command shows that the rlogin service has no dependents. The second command in this example disables the rlogin service. The third command shows that the state of the rlogin service instance is disabled.

```
svcs -D network/login:rlogin
svcadm disable network/login:rlogin
STATE STIME FMRI
svcs network/login:rlogin
STATE STIME FMRI
disabled 11:17:24 svc:/network/login:rlogin
```

## ▼ How to Enable a Service Instance

Use the following procedure to enable a service. The service status change is recorded in the service configuration repository. Once the service is enabled, the enabled state will persist across system reboots if the service dependencies are met.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 Determine whether service dependencies are satisfied.

If the service is enabled, then the service dependencies are satisfied. If not, use `svcadm enable -r FMRI` to recursively enable all dependencies.

```
svcs -l FMRI|grep enabled
```

**3 Enable a service.**

```
svcadm enable FMRI
```

**Example 19–8** Enabling the rlogin Service

The second command in this example enables the rlogin service. The third command shows that the state of the rlogin service instance is online.

```
svcs -l network/login:rlogin|grep enabled
enabled false
svcadm enable network/login:rlogin
svcs network/login:rlogin
STATE STIME FMRI
online 12:09:16 svc:/network/login:rlogin
```

**Example 19–9** Enabling a Service in Single-user Mode

The following command enables rpcbind. The -t option starts the service in temporary mode which does not change the service repository. The repository is not writable in single-user mode. The -r option recursively starts all the dependencies of the named service.

```
svcadm enable -rt rpc/bind
```

**▼ How to Restart a Service**

If a service is currently running but needs to be restarted due to a configuration change or some other reason, the service can be restarted without you having to type separate commands to stop and start the service. The only reason to specifically disable and then enable a service is if changes need to be made before the service is enabled, and after the service is disabled.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Restart a service.**

```
svcadm restart FMRI
```

**▼ How to Restore a Service That Is in the Maintenance State****1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Determine if any process that are dependent to the service have not stopped.**

Normally, when a service instance is in a maintenance state, all processes associated with that instance have stopped. However, you should make sure before you proceed. The following command lists all of the processes that are associated with a service instance as well as the PIDs for those processes.

```
svcs -p FMRI
```

**3 (Optional) Kill any remaining processes.**

Repeat this step for all processes that are displayed by the `svcs` command.

```
kill -9 PID
```

**4 If necessary, repair the service configuration.**

Consult the appropriate service log files in `/var/svc/log` for a list of errors.

**5 Restore the service.**

```
svcadm clear FMRI
```

## ▼ How to Revert to Another SMF Snapshot

If the service configuration is wrong, the problem can be fixed by reverting to the last snapshot that started successfully. In this procedure, a previous snapshot of the `console-login` service is used.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Run the `svccfg` command.**

```
svccfg
svc:>
```

**a. Select the service instance that you want to fix.**

---

**Note** – You must use an FMRI that fully defines the instance. No shortcuts are allowed.

---

```
svc:> select system/console-login:default
svc:/system/console-login:default>
```

**b. Generate a list of available snapshots.**

```
svc:/system/console-login:default> listsnap
initial
running
```



```
start
svc:/system/console-login:default>
```

### c. Select to revert to the start snapshot.

The start snapshot is the last snapshot in which the service successfully started.

```
svc:/system/console-login:default> revert start
svc:/system/console-login:default>
```

### d. Quit svccfg.

```
svc:/system/console-login:default> quit
#
```

## 3 Update the information in the service configuration repository.

This step updates the repository with the configuration information from the start snapshot.

```
svcadm refresh system/console-login
```

## 4 Restart the service instance.

```
svcadm restart system/console-login
```

## ▼ How to Create an SMF Profile

A profile is an XML file which lists SMF services and whether each should be enabled or disabled. Profiles are used to enable or disable many services at once. Not all services need to be listed in a profile. Each profile only needs to include those services that need to be enabled or disabled to make the profile useful.

### 1 Create a profile.

In this example, the svccfg command is used to create a profile which reflects which services are enabled or disabled on the current system. Alternately, you could make a copy of an existing profile to edit.

```
svccfg extract> profile.xml
```

If you are using Oracle Solaris JumpStart, if you have large numbers of identical systems, or if you want to archive the system configuration for later restoration, you may want to use this procedure to create a unique version of a SMF profile.

### 2 Edit the profile.xml file to make any required changes.

#### a. Change the name of the profile in the service\_bundle declaration.

In this example the name is changed to profile.

```
cat profile.xml
...
<service_bundle type='profile' name='profile'
 xmlns:xi='http://www.w3.org/2003/XInclude'
 ...
```

**b. Remove any services that should not be managed by this profile.**

For each service, remove the three lines that describe the service. Each service description starts with `<service` and ends with `</service>`. This example shows the lines for the LDAP client service.

```
cat profile.xml
...
<service name='network/ldap/client' version='1' type='service'>
 <instance name='default' enabled='true' />
</service>
```

**c. Add any services that should be managed by this profile.**

Each service needs to be defined using the three line syntax shown above.

**d. If necessary, change the enabled flag for selected services.**

In this example, the `sendmail` service is disabled.

```
cat profile.xml
...
<service name='network/smtp' version='1' type='service'>
 <instance name='sendmail' enabled='false' />
</service>
...
```

**3 When necessary, apply the new profile.**

See [“How to Apply an SMF Profile” on page 370](#) for instructions.

## ▼ How to Apply an SMF Profile

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in \*System Administration Guide: Security Services\*](#).

**2 Apply an profile.**

In this example, the `profile.xml` profile is used.

```
svccfg apply profile.xml
```

---

**Note** – For specific instructions for switching between the `generic_limited_net.xml` and `generic_open.xml` and the properties that need to be applied when making this switch, please see [“Changing Services Offered to the Network with `generic\*.xml`” on page 371](#)

---

## ▼ Changing Services Offered to the Network with `generic*.xml`

The `netservices` command switches system services between minimal network exposure and the traditional network exposure (as in previous Solaris releases). The switch is done with the `generic_limited.xml` and `generic_open.xml` profiles. In addition, some services properties are changed by the command to limit some services to a local-only mode or to the traditional mode, as appropriate.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 Run the `netservices` command.

In this example, the open or traditional network exposure is selected.

```
/usr/sbin/netservices open
```

#### Example 19–10 Limiting Network Service Exposure

This command changes properties to run some services in local mode, as well as restricts which services are enabled with the `generic_limited_net` profile. The command should only be used if the `generic_open.xml` profile had been applied.

```
/usr/sbin/netservices limited
```

## Configuring SMF Services

### ▼ How to Modify a Service

The following procedure shows how to change the configuration of a service that is not managed by the `inetd` service.

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

#### 2 Make changes to the configuration files, as needed.

Many of the services have one or more configuration files that are used to define the startup or other configuration information. These files can be changed while the service is running. The contents of the files is only checked when the service is started.

**3 Restart the service.**

```
svcadm restart FMRI
```

**Example 19–11 Sharing an NFS File System**

To share a file system using the NFS service, you must define the file system in the `/etc/dfs/dfstab` file and then restart the NFS service. This example shows you what the `dfstab` file could look like, as well as how to restart the service.

```
cat /etc/dfs/dfstab
:
:
share -F nfs -o rw /export/home
svcadm restart svc:/network/nfs/server
```

**▼ How to Change an Environment Variable for a Service**

This procedure shows how to modify `cron` environment variables to help with debugging.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

**2 Verify that the service is running.**

```
svcs system/cron
STATE STIME FMRI
online Dec_04 svc:/system/cron:default
```

**3 Set environment variables.**

In this example the `UMEM_DEBUG` and `LD_PRELOAD` environment variables are set. For information about the `setenv` subcommand refer to the `svccfg(1M)` man page.

```
svccfg -s system/cron:default setenv UMEM_DEBUG default
svccfg -s system/cron:default setenv LD_PRELOAD libumem.so
```

**4 Refresh and restart the service.**

```
svcadm refresh system/cron
svcadm restart system/cron
```

**5 Verify that the change has been made.**

```
pargs -e 'pgrep -f /usr/sbin/cron'
100657: /usr/sbin/cron
envp[0]: LOGNAME=root
envp[1]: LD_PRELOAD=libumem.so
envp[2]: PATH=/usr/sbin:/usr/bin
envp[3]: SMF_FMRI=svc:/system/cron:default
envp[4]: SMF_METHOD=/lib/svc/method/svc-cron
```

```

envp[5]: SMF_RESTARTER=svc:/system/svc/restarter:default
envp[6]: TZ=GB
envp[7]: UMEM_DEBUG=default
#

```

## ▼ How to Change a Property for an inetd Controlled Service

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 List the properties for the specific service.

This command displays all of the properties for the service identified by the FMRI.

```
inetadm -l FMRI
```

### 3 Change the property for the service.

Each property for an inetd controlled service is defined by a property name and an assigned value. Supplying the property name without a specified value resets the property to the default value. Specific information about the properties for a service should be covered in the man page associated with the service.

```
inetadm -m FMRI property-name=value
```

### 4 Verify that the property has changed.

List the properties again to make sure that the appropriate change has occurred.

```
inetadm -l FMRI
```

### 5 Confirm that the change has taken effect.

Confirm the property change that the change has the desired effect.

#### Example 19–12 Changing the tcp\_trace Property for telnet

The following example shows how to set the tcp\_trace property for telnet to true. Checking the syslog output after running a telnet command shows that the change has taken effect.

```

inetadm -l svc:/network/telnet:default
SCOPE NAME=VALUE
 name="telnet"
.
.
default inherit_env=TRUE
default tcp_trace=FALSE
default tcp_wrappers=FALSE
inetadm -m svc:/network/telnet:default tcp_trace=TRUE

```

```

inetadm -l svc:/network/telnet:default
SCOPE NAME=VALUE
 name="telnet"
.
.
default inherit_env=TRUE
 tcp_trace=TRUE
default tcp_wrappers=FALSE
telnet localhost
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.
login: root
Password:
Last login: Mon Jun 21 05:55:45 on console
Oracle Corporation SunOS 5.10 Generic Patch January 2005
^D
Connection to localhost closed by foreign host.
tail -1 /var/adm/messages
Jun 21 06:04:57 yellow-19 inetd[100308]: [ID 317013 daemon.notice] telnet[100625]
 from 127.0.0.1 32802

```

## ▼ How to Modify a Command-Line Argument for an inetd Controlled Service

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 List the exec property for the specific service.

This command displays all the properties for the service identified by the FMRI. Adding the `grep` command restricts the output to the `exec` property for the service.

```
inetadm -l FMRI|grep exec
```

### 3 Change the exec property for the service.

The *command-syntax* set with the `exec` property defines the command string that is run when the service is started.

```
inetadm -m FMRI exec="command-syntax"
"
```

### 4 Verify that the property has changed.

List the properties again to make sure that the appropriate change has occurred.

```
inetadm -l FMRI
```

**Example 19-13** Adding the Connection Logging (-l) Option to the ftp Command

In this example, the -l option is added to the ftp daemon when it is started. The effect of this change can be seen by reviewing the syslog output after a ftp login session has been completed.

```
inetadm -l svc:/network/ftp:default | grep exec
 exec="/usr/sbin/in.ftpd -a"
inetadm -m svc:/network/ftp:default exec="/usr/sbin/in.ftpd -a -l"
inetadm -l svc:/network/ftp:default
SCOPE NAME=VALUE
 name="ftp"
 endpoint_type="stream"
 proto="tcp6"
 isrpc=FALSE
 wait=FALSE
 exec="/usr/sbin/in.ftpd -a -l"
.
.
ftp localhost
Connected to localhost.
220 yellow-19 FTP server ready.
Name (localhost:root): mylogin
331 Password required for mylogin.
Password:
230 User mylogin logged in.
Remote system type is UNIX.
Using binary mode to transfer files.
ftp> quit
221-You have transferred 0 bytes in 0 files.
221-Total traffic for this session was 236 bytes in 0 transfers.
221-Thank you for using the FTP service on yellow-19.
221 Goodbye.
tail -2 /var/adm/messages
Jun 21 06:54:33 yellow-19 ftpd[100773]: [ID 124999 daemon.info] FTP LOGIN FROM localhost
[127.0.0.1], mylogin
Jun 21 06:54:38 yellow-19 ftpd[100773]: [ID 528697 daemon.info] FTP session closed
```

**▼ How to Convert inetd.conf Entries**

The following procedure converts `inetd.conf` entries into SMF service manifests. This procedure needs to be run any time a third-party application that depends on `inetd` is added to a system. Also run this procedure, if you need to make configuration changes to the entry in `/etc/inetd.conf`.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Convert the `inetd.conf` entries.**

The `inetconv` command converts each entry in the selected file into service manifests.

```
inetconv -i filename
```

**Example 19–14** Converting `/etc/inet/inetd.conf` Entries into SMF Service Manifests

```
inetconv -i /etc/inet/inetd.conf
```

## Using Run Control Scripts (Task Map)

| Task                          | Description                                                                       | For Instructions                                                                |
|-------------------------------|-----------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Stop or start a service.      | Use a run control script to stop or start a service.                              | “How to Use a Run Control Script to Stop or Start a Legacy Service” on page 376 |
| Add a run control script.     | Create a run control script and add it to the <code>/etc/init.d</code> directory. | “How to Add a Run Control Script” on page 377                                   |
| Disable a run control script. | Disable a run control script by renaming the file.                                | “How to Disable a Run Control Script” on page 378                               |

## Using Run Control Scripts

### ▼ How to Use a Run Control Script to Stop or Start a Legacy Service

One advantage of having individual scripts for each run level is that you can run scripts in the `/etc/init.d` directory individually to stop system services without changing a system's run level.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Stop the system service.**

```
/etc/init.d/filename
stop
```

**3 Restart the system service.**

```
/etc/init.d/filename
start
```



**4 Verify that the service has been stopped or started.**

```
pgrep -f service
```

**Example 19–15 Using a Run Control Script to Stop or Start a Service**

For example, you can stop the NFS server daemons by typing the following:

```
/etc/init.d/nfs.server stop
pgrep -f nfs
```

Then, you can restart the NFS server daemons by typing the following:

```
/etc/init.d/nfs.server start
pgrep -f nfs
101773
101750
102053
101748
101793
102114
pgrep -f nfs -d, | xargs ps -fp
 UID PID PPID C STIME TTY TIME CMD
daemon 101748 1 0 Sep 01 ? 0:06 /usr/lib/nfs/nfsmapid
daemon 101750 1 0 Sep 01 ? 26:27 /usr/lib/nfs/lockd
daemon 101773 1 0 Sep 01 ? 5:27 /usr/lib/nfs/statd
 root 101793 1 0 Sep 01 ? 19:42 /usr/lib/nfs/mountd
daemon 102053 1 0 Sep 01 ? 2270:37 /usr/lib/nfs/nfsd
daemon 102114 1 0 Sep 01 ? 0:35 /usr/lib/nfs/nfs4cbd
```

**▼ How to Add a Run Control Script**

If you want to add a run control script to start and stop a service, copy the script into the `/etc/init.d` directory. Then, create links in the `rcn.d` directory where you want the service to start and stop.

See the README file in each `/etc/rcn.d` directory for more information on naming run control scripts. The following procedure describes how to add a run control script.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Add the script to the `/etc/init.d` directory.**

```
cp filename/etc/init.d
chmod 0744 /etc/init.d/filename
chown root:sys /etc/init.d/filename
```

**3 Create links to the appropriate rc n.d directory.**

```
cd /etc/init.d
ln filename /etc/rc2.d/Snnfilename
ln filename /etc/rcn.d/Knnfilename
```

**4 Verify that the script has links in the specified directories.**

```
ls /etc/init.d/*filename /etc/rc2.d/*filename /etc/rcn.d/*filename
```

**Example 19–16 Adding a Run Control Script**

The following example shows how to add a run control script for the xyz service.

```
cp xyz /etc/init.d
chmod 0744 /etc/init.d/xyz
chown root:sys /etc/init.d/xyz
cd /etc/init.d
ln xyz /etc/rc2.d/S99xyz
ln xyz /etc/rc0.d/K99xyz
ls /etc/init.d/*xyz /etc/rc2.d/*xyz /etc/rc0.d/*xyz
```

**▼ How to Disable a Run Control Script**

You can disable a run control script by renaming it with an underscore ( `_` ) at the beginning of the file name. Files that begin with an underscore or dot are not executed. If you copy a file by adding a suffix to it, both files will be run.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

**2 Rename the script by adding an underscore ( `_` ) to the beginning of the new file.**

```
cd /etc/rcn.d
mv filename_filename
```

**3 Verify that the script has been renamed.**

```
ls _*
_filename
```

**Example 19–17 Disabling a Run Control Script**

The following example shows how to rename the S99datainit script.

```
cd /etc/rc2.d
mv S99datainit _S99datainit
ls _*
_S99datainit
```

# Troubleshooting the Service Management Facility

## ▼ Debugging a Service That Is Not Starting

In this procedure, the print service is disabled.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

### 2 Request information about the hung service.

```
svcs -xv
svc:/application/print/server:default (LP Print Service)
State: disabled since Wed 13 Oct 2004 02:20:37 PM PDT
Reason: Disabled by an administrator.
See: http://sun.com/msg/SMF-8000-05
See: man -M /usr/share/man -s 1M lpsched
Impact: 2 services are not running:
 svc:/application/print/rfc1179:default
 svc:/application/print/ipp-listener:default
```

The -x option provides additional information about the service instances that are impacted.

### 3 Enable the service.

```
svcadm enable application/print/server
```

## ▼ How to Repair a Corrupt Repository

This procedure shows how to replace a corrupt repository with a default copy of the repository. When the repository daemon, `svc.configd`, is started, it does an integrity check of the configuration repository. This repository is stored in `/etc/svc/repository.db`. The repository can become corrupted due to one of the following reasons:

- Disk failure
- Hardware bug
- Software bug
- Accidental overwrite of the file

If the integrity check fails, the `svc.configd` daemon writes a message to the console similar to the following:

```
svc.configd: smf(5) database integrity check of:
 /etc/svc/repository.db
```

failed. The database might be damaged or a media error might have prevented it from being verified. Additional information useful to your service provider is in:

```
/etc/svc/volatile/db_errors
```

The system will not be able to boot until you have restored a working database. `svc.startd(1M)` will provide a `sulogin(1M)` prompt for recovery purposes. The command:

```
/lib/svc/bin/restore_repository
```

can be run to restore a backup version of your repository. See <http://sun.com/msg/SMF-8000-MY> for more information.

The `svc.startd` daemon then exits and starts `sulogin` to enable you to perform maintenance.

- 1 **Enter the root password at the `suLogin` prompt. `suLogin` enables the root user to enter system maintenance mode to repair the system.**
- 2 **Run the following command:**

```
/lib/svc/bin/restore_repository
```

Running this command takes you through the necessary steps to restore a non-corrupt backup. SMF automatically takes backups of the repository at key system moments. For more information see “[SMF Repository Backups](#)” on page 353.

When started, the `/lib/svc/bin/restore_repository` command displays a message similar to the following:

```
Repository Restore utility
See http://sun.com/msg/SMF-8000-MY for more information on the use of
this script to restore backup copies of the smf(5) repository.
```

```
If there are any problems which need human intervention, this script
will give instructions and then exit back to your shell.
```

```
Note that upon full completion of this script, the system will be
rebooted using reboot(1M), which will interrupt any active services.
```

```
If the system that you are recovering is not a local zone, the script explains how to remount the /
and /usr file systems with read and write permissions to recover the databases. The script exits
after printing these instructions. Follow the instructions, paying special attention to any errors
that might occur.
```

```
After the root (/) file system is mounted with write permissions, or if the system is a local zone,
you are prompted to select the repository backup to restore:
```

```
The following backups of /etc/svc/repository.db exists, from
oldest to newest:
```

```
... list of backups ...
```

Backups are given names, based on type and the time the backup was taken. Backups beginning with `boot` are completed before the first change is made to the repository after system boot. Backups beginning with `manifest_import` are completed after `svc:/system/manifest-import:default` finishes its process. The time of the backup is given in `YYYYMMDD_HHMMSS` format.

### 3 Enter the appropriate response.

Typically, the most recent backup option is selected.

Please enter one of:

- 1) `boot`, for the most recent post-boot backup
- 2) `manifest_import`, for the most recent `manifest_import` backup.
- 3) a specific backup repository from the above list
- 4) `-seed-`, the initial starting repository. (All customizations will be lost.)
- 5) `-quit-`, to cancel.

Enter response [boot]:

If you press `Enter` without specifying a backup to restore, the default response, enclosed in `[]` is selected. Selecting `-quit-` exits the `restore_repository` script, returning you to your shell prompt.

---

**Note** – Selecting `-seed-` restores the seed repository. This repository is designed for use during initial installation and upgrades. Using the seed repository for recovery purposes should be a last resort.

---

After the backup to restore has been selected, it is validated and its integrity is checked. If there are any problems, the `restore_repository` command prints error messages and prompts you for another selection. Once a valid backup is selected, the following information is printed, and you are prompted for final confirmation.

After confirmation, the following steps will be taken:

```
svc.startd(1M) and svc.configd(1M) will be quiesced, if running.
/etc/svc/repository.db
-- renamed --> /etc/svc/repository.db_old_YYYYMMDD_HHMMSS
/etc/svc/volatile/db_errors
-- copied --> /etc/svc/repository.db_old_YYYYMMDD_HHMMSS_errors
repository_to_restore
-- copied --> /etc/svc/repository.db
and the system will be rebooted with reboot(1M).
```

Proceed [yes/no]?

### 4 Type yes to remedy the fault.

The system reboots after the `restore_repository` command executes all of the listed actions.

## ▼ How to Boot Without Starting Any Services

If problems with starting services occur, sometimes a system will hang during the boot. This procedure shows how to troubleshoot this problem.

### 1 Boot without starting any services.

This command instructs the `svc.startd` daemon to temporarily disable all services and start `su` login on the console.

```
ok boot -m milestone=none
```

### 2 Log in to the system as root.

### 3 Enable all services.

```
svcadm milestone all
```

### 4 Determine where the boot process is hanging.

When the boot process hangs, determine which services are not running by running `svcs -a`. Look for error messages in the log files in `/var/svc/log`.

### 5 After fixing the problems, verify that all services have started.

#### a. Verify that all needed services are online.

```
svcs -x
```

#### b. Verify that the `console-login` service dependencies are satisfied.

This command verifies that the `login` process on the console will run.

```
svcs -l system/console-login:default
```

### 6 Continue the normal booting process.

## ▼ How to Force a `su` login Prompt If the `system/filesystem/local:default` Service Fails During Boot

Local file systems that are not required to boot the Solaris OS are mounted by the `svc:/system/filesystem/local:default` service. When any of those file systems are unable to be mounted, the service enters a maintenance state. System startup continues, and any services which do not depend on `filesystem/local` are started. Services which require `filesystem/local` to be online before starting through dependencies are not started.

To change the configuration of the system so that a `su` login prompt appears immediately after the service fails instead of allowing system startup to continue, follow the procedure below.

## 1 Modify the system/console-login service.

```
svccfg -s svc:/system/console-login
svc:/system/console-login> addpg site,filesystem-local dependency
svc:/system/console-login> setprop site,filesystem-local/entities = fmri: svc:/system/filesystem/local

svc:/system/console-login> setprop site,filesystem-local/grouping = astring: require_all

svc:/system/console-login> setprop site,filesystem-local/restart_on = astring: none

svc:/system/console-login> setprop site,filesystem-local/type = astring: service

svc:/system/console-login> end
```

## 2 Refresh the service.

```
svcadm refresh console-login
```

### Example 19–18 Forcing an sulogin Prompt Using Oracle Solaris JumpStart

Save the following commands into a script and save it as `/etc/rcS.d/S01site-customfs`.

```
#!/bin/sh
#
This script adds a dependency from console-login -> filesystem/local
This forces the system to stop the boot process and drop to an sulogin prompt
if any file system in filesystem/local fails to mount.

PATH=/usr/sbin:/usr/bin
export PATH

 svccfg -s svc:/system/console-login << EOF
addpg site,filesystem-local dependency
setprop site,filesystem-local/entities = fmri: svc:/system/filesystem/local
setprop site,filesystem-local/grouping = astring: require_all
setprop site,filesystem-local/restart_on = astring: none
setprop site,filesystem-local/type = astring: service
EOF

svcadm refresh svc:/system/console-login

[-f /etc/rcS.d/S01site-customfs] &&
 rm -f /etc/rcS.d/S01site-customfs
```

**Troubleshooting** When a failure occurs with the `system/filesystem/local:default` service, the `svcs -vx` command should be used to identify the failure. After the failure has been fixed, the following command clears the error state and allows the system boot to continue: `svcadm clear filesystem/local`.





## Managing Software (Overview)

---

Software management involves adding and removing software from stand-alone systems, servers, and their clients. This chapter describes the various tools that are available for installing and managing software.

This chapter does not describe installing the Oracle Solaris OS on a new system, nor does it describe installing or upgrading to a new version of the Oracle Solaris OS. For information about installing or upgrading the Oracle Solaris OS, see *Oracle Solaris 10 8/11 Installation Guide: Basic Installations*.

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

- “What's New in Software Management in the Oracle Solaris Operating System?” on page 386
- “Where to Find Software Management Tasks” on page 389
- “Overview of Software Packages” on page 389
- “Tools for Managing Software Packages” on page 390
- “Adding or Removing a Software Package (pkgadd)” on page 391
- “Key Points for Adding Software Packages (pkgadd)” on page 391
- “Guidelines for Removing Packages (pkgrm)” on page 392
- “Restrictions on Adding and Removing Software Packages and Patches for Solaris Releases That Are Not Zones Aware” on page 392
- “Avoiding User Interaction When Adding Packages (pkgadd)” on page 393

For step-by-step instructions on managing software, see Chapter 21, “Managing Software With Oracle Solaris System Administration Tools (Tasks),” and Chapter 22, “Managing Software by Using Oracle Solaris Package Commands (Tasks).”

For information about managing software on systems that have Oracle Solaris zones installed, see Chapter 26, “Adding and Removing Packages and Patches on a Solaris System With Zones Installed (Tasks),” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*.

# What's New in Software Management in the Oracle Solaris Operating System?

This section describes the new software management features in this Oracle Solaris release.

For a complete listing of new features and a description of Oracle Solaris releases, see the [Oracle Solaris 10 8/11 What's New](#).

## Oracle Solaris Auto Registration

**Oracle Solaris 10 9/10:** Auto Registration is a built-in mechanism that enables configuration data about your system to be gathered and automatically registered with the Oracle Product Registration System for inventory and support purposes. The `regadm` command can be used by a privileged system administrator to manage the Auto Registration feature and to explicitly register a system, independent of the installation or upgrade process.

For more information, see [Chapter 17, “Working With the Oracle Solaris Auto Registration `regadm` Command \(Tasks\)”](#).

## Package and Patching Tools Enhancements to Support Oracle Solaris Zones

**Oracle Solaris 10 10/09:** Starting with this release, the following enhancements have been made to the package and patching tools to support systems that have multiple non-global zones installed:

- **Parallel Patching of Zones**

The Parallel Patching of Zones feature introduces enhancements to the standard Oracle Solaris 10 patching tools to enable patching of non-global zones in parallel. This feature enhancement improves patching performance when multiple non-global zones are installed on a system by accelerating the time it takes to patch systems. Because zones are isolated environments, it is safe to patch multiple zones in parallel. This ability applies to both sparse zones and whole root zones.

These improvements primarily include changes to the `patchadd` and `patchrm` commands. A new configuration file, `/etc/patch/pdo.conf`, can be modified to specify a number that defines how many zones are patched in parallel.

---

**Note** – If you are not running at least the Solaris 10 10/09 release, the feature is implemented through patch 119254-66, or later revision (SPARC), and 119255-66, or later revision (x86). The global zone is still patched before the non-global zones are patched.

---

For more information, see the [patchadd\(1M\)](#) and the [patchrm\(1M\)](#) man pages.

- **Turbo-Charging of SVr4 Packaging Commands**

This feature accelerates the time it takes to install the Solaris OS, use Solaris Live Upgrade, or install non-global zones by using SVr4 package commands. Note that this feature cannot be disabled.

For more details about these and other enhancements that support the use of package and patch tools in a zones environment, see the following documentation:

- Chapter 26, “Adding and Removing Packages and Patches on a Solaris System With Zones Installed (Tasks),” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*
- Chapter 23, “Moving and Migrating Non-Global Zones (Tasks),” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*
- Chapter 29, “Upgrading a Solaris 10 System That Has Installed Non-Global Zones,” in *System Administration Guide: Oracle Solaris Containers-Resource Management and Oracle Solaris Zones*

## Deferred-Activation Patching

Patching tools have changed to handle larger patches. Starting with patch 119254-42 and 119255-42, the patch installation commands, `patchadd` and `patchrm`, have been modified to change the way in which certain patches that deliver new features are handled. This modification affects the installation of these patches on any Oracle Solaris OS. These *deferred-activation* patches are better equipped to handle the large scope of changes that are delivered in feature patches.

For more details, see <http://www.oracle.com/technetwork/systems/index.html>.

## Common Agent Container Included in the Oracle Solaris OS

The Common Agent Container is a stand-alone Java program that implements a container for Java management applications. This program provides a management infrastructure that is designed for Java Management Extensions (JMX) and Java Dynamic Management Kit (Java DMK) based management functionality. The software is installed by the `SUNWcacao-rt` package and resides in the `/usr/lib/cacao` directory.

Typically, the container is not visible.

However, there are two instances when you might need to interact with the container daemon:

- In the event that another application attempts to use a network port that is reserved for the Common Agent Container.
- In the event that a certificate store is compromised. If this conflict occurs, you might have to regenerate the Common Agent Container certificate keys.

For information about how to troubleshoot these problems, see “[Troubleshooting Common Agent Container Problems in the Oracle Solaris OS](#)” in *System Administration Guide: Advanced Administration*.

## Improvements to How `patchadd -M` Command Handles Multiple Patches

**Oracle Solaris 10:** Starting with this release, the functionality of the `patchadd -M` command is improved to more effectively and efficiently handle multiple patches and any dependencies between patches. As a result, you no longer have to specify patch IDs in numerical order when using the `patchadd` command.

Note that if you use the `patchadd -M` command without specifying a patch ID or patch IDs, all the patches in the directory are automatically installed on the system. To install a specific patch, or patches, you must specify the patch ID when using the `patchadd -M` command.

For more information, see the [patchadd\(1M\)](#) man page.

## Package and Patch Tool Enhancements

**Oracle Solaris 10:** The package and patch tools were enhanced in Oracle Solaris 10 to provide improved performance and extended functionality.

As a part of these enhancements, the `pkgchk` command now provides a new option to assist you in mapping files to packages. To map files to packages, use the `pkgchk -P` option instead of `grep pattern/var/sadm/install/contents`. The `-P` option enables you to use a partial path. Use this option with the `-l` option to list the information about the files that contain the partial path. For more information see “[How to Check the Integrity of Installed Objects \( `pkgchk -p`, `pkgchk -P` \)](#)” on [page 425](#) and the [pkgchk\(1M\)](#) man page.

## Where to Find Software Management Tasks

Use this table to find step-by-step instructions for managing software.

| Software Management Topics                              | For More Information                                                                                                                                                          |
|---------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Installing Oracle Solaris 10                            | <i>Oracle Solaris 10 8/11 Installation Guide: Basic Installations</i>                                                                                                         |
| Adding or removing software packages after installation | Chapter 21, “Managing Software With Oracle Solaris System Administration Tools (Tasks),” and Chapter 22, “Managing Software by Using Oracle Solaris Package Commands (Tasks)” |
| Adding or removing Solaris patches after installation   | Chapter 23, “Managing Patches”                                                                                                                                                |
| Troubleshooting software package problems               | Chapter 21, “Troubleshooting Software Package Problems (Tasks),” in <i>System Administration Guide: Advanced Administration</i>                                               |

## Overview of Software Packages

Software management involves installing or removing software products. Sun and its third-party ISVs deliver software as a collection of one or more *packages*.

The term *packaging* generically refers to the method for distributing and installing software products to systems where the products will be used. A package is a collection of files and directories in a defined format. This format conforms to the application binary interface (ABI), which is a supplement to the System V Interface Definition. The Solaris OS provides a set of utilities that interpret this format and provide the means to install a package, to remove a package, or to verify a package installation.

A *patch* is an accumulation of fixes for a known or potential problem within the Solaris OS or other supported software. A patch can also provide a new feature or an enhancement to a particular software release. A patch consists of files and directories that replace or update existing files and directories. Most Solaris patches are delivered as a set of sparse packages.

A *sparse package* contains only those objects that have been altered since the packages were first delivered as part of the Solaris distribution. Sparse packages accommodate patches that are smaller than if they were redistributed as complete packages to provide software updates. Delivering sparse packages also minimizes the changes that are made to the customer's environment. For more information about patches, see [Chapter 23, “Managing Patches.”](#)

## Tools for Managing Software Packages

The following table describes the tools for adding and removing software packages from a system after Oracle Solaris is installed on a system.

TABLE 20-1 Tools or Commands for Managing Software Packages

| Tool or Command                                                   | Description                                                                                                                                                                                                                                                                                 | Man Page                      |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| <code>installer</code>                                            | Launches an installer, such as the Oracle Solaris installation GUI, to add software from the Oracle Solaris installation media. The installer must be available either locally or remotely.                                                                                                 | <a href="#">installer(1M)</a> |
| <code>prodreg (GUI)</code>                                        | Launches an installer to add, remove, or display software product information. Use Oracle Solaris Product Registry to remove or display information about software products that were originally installed by using the Oracle Solaris installation GUI or the <code>pkgadd</code> command. | <a href="#">prodreg(1M)</a>   |
| Oracle Solaris Product Registry <code>prodreg Viewer (CLI)</code> | Use the <code>prodreg</code> command to remove or display information about software products that were originally installed by using the Oracle Solaris installation GUI or the <code>pkgadd</code> command.                                                                               | <a href="#">prodreg(1M)</a>   |
| <code>pkgadd</code>                                               | Installs software package.                                                                                                                                                                                                                                                                  | <a href="#">pkgadd(1M)</a>    |
| <code>pkgchk</code>                                               | Checks the installation of a software package.                                                                                                                                                                                                                                              | <a href="#">pkgchk(1M)</a>    |
| <code>pkginfo</code>                                              | Lists software package information.                                                                                                                                                                                                                                                         | <a href="#">pkginfo(1)</a>    |
| <code>pkgparam</code>                                             | Displays software package parameter values.                                                                                                                                                                                                                                                 | <a href="#">pkgparam(1)</a>   |
| <code>pkgrm</code>                                                | Removes a software package.                                                                                                                                                                                                                                                                 | <a href="#">pkgrm(1M)</a>     |
| <code>pkgtrans</code>                                             | Translates an installable package from one format to another format. The <code>-g</code> option instructs the <code>pkgtrans</code> command to generate and store a signature in the resulting data stream.                                                                                 | <a href="#">pkgtrans(1)</a>   |

For more information about these commands, see [Chapter 21, “Managing Software With Oracle Solaris System Administration Tools \(Tasks\)”](#) and [Chapter 22, “Managing Software by Using Oracle Solaris Package Commands \(Tasks\)”](#).

## Adding or Removing a Software Package (pkgadd)

All the software management tools that are listed in [Table 20–1](#) are used to add, remove, or query information about installed software. The Solaris Product Registry `prodreg` viewer and the Solaris installation GUI both access install data that is stored in the Solaris Product Registry. The package tools, such as the `pkgadd` and `pkgrm` commands, also access or modify install data.

When you add a package, the `pkgadd` command uncompresses and copies files from the installation media to a local system's disk. When you remove a package, the `pkgrm` command deletes all files associated with that package, unless those files are also shared with other packages.

Package files are delivered in package format and are unusable as they are delivered. The `pkgadd` command interprets the software package's control files, and then uncompresses and installs the product files onto the system's local disk.

Although the `pkgadd` and `pkgrm` commands do not log their output to a standard location, they do keep track of the package that is installed or removed. The `pkgadd` and `pkgrm` commands store information about a package that has been installed or removed in a software product database.

By updating this database, the `pkgadd` and `pkgrm` commands keep a record of all software products installed on the system.

## Key Points for Adding Software Packages (pkgadd)

Keep the following key points in mind before you install or remove packages on your system:

- **Package naming conventions** – Sun packages always begin with the prefix `SUNW`, as in `SUNWaccr`, `SUNWadmap`, and `SUNWcsu`. Third-party packages usually begin with a prefix that corresponds to the company's stock symbol.
- **What software is already installed** – You can use the Solaris installation GUI, Solaris Product Registry `prodreg` viewer (either GUI or CLI) or the `pkginfo` command to determine the software that is already installed on a system.
- **How servers and clients share software** – Clients might have software that resides partially on a server and partially on the client. In such cases, adding software for the client requires that you add packages to both the server and the client.

## Guidelines for Removing Packages (pkgrm)

Even though you might be tempted to use the `rm` command to remove a package, you should use one of the tools that is listed in [Table 20–1](#) instead. For example, you could use the `rm` command to remove a binary executable file. However, doing so is not the same as using the `pkgrm` command to remove the software package that includes that binary executable. Using the `rm` command to remove a package's files will corrupt the software products database. If you really only want to remove one file, you can use the `removef` command. This command will update the software product database correctly so that the file is no longer a part of the package. For more information, see the [removef\(1M\)](#) man page.

If you intend to keep multiple versions of a package, install new versions into a different directory than the already installed package by using the `pkgadd` command. For example, if you intended to keep multiple versions of a document processing application. The directory where a package is installed is referred to as the base directory. You can manipulate the base directory by setting the `basedir` keyword in a special file called an administration file. For more information on using an *administration file* and on setting the base directory, see “[Avoiding User Interaction When Adding Packages \(pkgadd\)](#)” on page 393 and the `admin(4)` man page.

---

**Note** – If you use the upgrade option when installing Solaris software, the Solaris installation software checks the software product database to determine the products that are already installed on the system.

---

## Restrictions on Adding and Removing Software Packages and Patches for Solaris Releases That Are Not Zones Aware

On systems that are running an Oracle Solaris release that is not zones aware, using any command that accepts the `-R` option to specify an alternate root path for a global zone that has non-global zones installed, does not work.

These commands include:

- `pkgadd`
- `pkgrm`
- `patchadd`
- `patchrm`

See the [pkgadd\(1M\)](#), [pkgrm\(1M\)](#), [patchadd\(1M\)](#), and [patchrm\(1M\)](#) man pages.



# Avoiding User Interaction When Adding Packages (pkgadd)

This section provides information about avoiding user interaction when adding packages with the `pkgadd` command.

## Using an Administration File

When you use the `pkgadd -a` command, the command consults a special administration file for information about how the installation should proceed. Normally, the `pkgadd` command performs several checks and prompts the user for confirmation before it actually adds the specified package. You can, however, create an administration file that indicates to the `pkgadd` command that it should bypass these checks and install the package without user confirmation.

The `pkgadd` command, by default, checks the current working directory for an administration file. If the `pkgadd` command doesn't find an administration file in the current working directory, it checks the `/var/sadm/install/admin` directory for the specified administration file. The `pkgadd` command also accepts an absolute path to the administration file.

---

**Note** – Use administration files judiciously. You should know where a package's files are installed and how a package's installation scripts run before using an administration file to avoid the checks and prompts that the `pkgadd` command normally provides.

---

The following example shows an administration file that prevents the `pkgadd` command from prompting the user for confirmation before installing the package.

```
mail=
instance=overwrite
partial=nocheck
runlevel=nocheck
idepend=nocheck
rdepend=nocheck
space=nocheck
setuid=nocheck
conflict=nocheck
action=nocheck
networktimeout=60
networkretries=3
authentication=quit
keystore=/var/sadm/security
proxy=
basedir=default
```

Besides using administration files to avoid user interaction when you add packages, you can use them in several other ways. For example, you can use an administration file to quit a package installation (without user interaction) if there's an error or to avoid interaction when you remove packages by using the `pkgrm` command.

You can also assign a special installation directory for a package, which you might do if you wanted to maintain multiple versions of a package on a system. To do so, set an alternate base directory in the administration file by using the `basedir` keyword. The keyword specifies where the package will be installed. For more information, see the [admin\(4\)](#) man page.

## Using a Response File (pkgadd)

A response file contains your answers to specific questions that are asked by an *interactive package*. An interactive package includes a request script that asks you questions prior to package installation, such as whether optional pieces of the package should be installed.

If you know prior to installation that the package is an interactive package, and you want to store your answers to prevent user interaction during future installations, use the `pkgask` command to save your response. For more information on this command, see [pkgask\(1M\)](#).

Once you have stored your responses to the questions asked by the request script, you can use the `pkgadd -r` command to install the package without user interaction.

# Managing Software With Oracle Solaris System Administration Tools (Tasks)

---

This chapter describes how to add, verify, and remove software packages by using the Oracle Solaris installation graphical user interface (GUI) and the Oracle Solaris Product Registry.

For information about software management features that are new in this release, see [“What's New in Software Management in the Oracle Solaris Operating System?”](#) on page 386.

For information about the procedures that are associated with performing software management tasks, see:

- [“Adding Software With the Oracle Solaris Installation GUI”](#) on page 396
- [“Managing Software With the Oracle Solaris Product Registry GUI \(Task Map\)”](#) on page 397
- [“Managing Software With the Oracle Solaris Product Registry Command-Line Interface \(Task Map\)”](#) on page 401

## Oracle Solaris Product Registry and Solaris GUI Installation Tools for Managing Software

The following table lists the commands to use for adding, removing, and checking the installation of software packages the Oracle Solaris installation GUI and Oracle Solaris Product Registry tools.

TABLE 21-1 System Administration Tools for Managing Software Packages

| Tool                   | Description                                                                                      | Man Page                      |
|------------------------|--------------------------------------------------------------------------------------------------|-------------------------------|
| <code>installer</code> | Installs or removes a software package with an installer                                         | <a href="#">installer(1M)</a> |
| <code>prodreg</code>   | Enables you to browse, unregister, and uninstall software in the Oracle Solaris Product Registry | <a href="#">prodreg(1M)</a>   |

## Adding Software With the Oracle Solaris Installation GUI

This section describes how to use the Oracle Solaris installation GUI to add software to a system on which you have installed the Oracle Solaris OS. The installation GUI installs only the components of the software groups that you skipped when you initially installed Oracle Solaris. You cannot upgrade to another software group after installing or upgrading the OS.

### ▼ How to Install Software With the Oracle Solaris Installation GUI Program

---

**Note** – This procedure assumes that the system is running volume management (`vol`). If your system is not running volume management, see [Chapter 3, “Accessing Removable Media \(Tasks\)”](#) in *System Administration Guide: Devices and File Systems*. This chapter provides information about accessing removable media without volume management.

---

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.

#### 2 Depending on whether you are installing the OS from a CD, DVD, or from the network, do one of the following:

- If you are running a release that uses CDs, insert the CD into the CD-ROM drive.  
If you insert the Solaris 10 Languages CD, the installation GUI starts automatically. Proceed to [Step 5](#).

---

**Note** – In this Oracle Solaris release, the software is delivered on DVD *only*.

---

- If you are installing from a DVD, insert the DVD into the DVD-ROM drive.
- If you are installing from the network, locate the net image of the software that you want to install.

#### 3 Change directories to find the Oracle Solaris installation GUI installer.

Oracle Solaris installation GUI installers are located in various directories on the CDs and on the DVD.

- Oracle Solaris 10 Software CDs or DVD.
- Oracle Solaris 10 Documentation DVD.

- Oracle Solaris 10 Languages CD. The Solaris installation GUI starts automatically when the CD is inserted.
- 4 **Follow the instructions to install the software.**
    - From the command line, type the following command:
 

```
% ./installer [options]
```

      - nodisplay   Runs the installer without a GUI.
      - noconsole   Runs without any interactive text console device. Use this option with the -nodisplay option when you include the installer command in a UNIX script for installing software.
    - From a file manager, double-click Installer or installer.  
An Installer window is displayed, followed by the installation GUI dialog box.
  - 5 **Follow the directions on the screen to install the software.**
  - 6 **When you have finished adding software, click Exit.**  
The installation GUI exits.

## Managing Software With the Oracle Solaris Product Registry GUI (Task Map)

The following task map describes the software management tasks that you can perform with the Oracle Solaris Product Registry.

| Task                                                                             | Description                                                                                                                                                                        | For Instructions                                                                                                                     |
|----------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|
| View installed or uninstalled software with the Oracle Solaris Product Registry. | Used for learning about installed or uninstalled software.                                                                                                                         | <a href="#">“How to View Installed or Uninstalled Software Information With the Oracle Solaris Product Registry GUI” on page 398</a> |
| Install software with the Oracle Solaris Product Registry.                       | You can use the Oracle Solaris Product Registry to find software and launch the Oracle Solaris installation GUI. This program takes you through the installation of that software. | <a href="#">“How to Install Software With the Oracle Solaris Product Registry GUI” on page 399</a>                                   |

| Task                                                         | Description                                                         | For Instructions                                                                                     |
|--------------------------------------------------------------|---------------------------------------------------------------------|------------------------------------------------------------------------------------------------------|
| Uninstall software with the Oracle Solaris Product Registry. | Use to uninstall software with the Oracle Solaris Product Registry. | <a href="#">“How to Uninstall Software With the Oracle Solaris Product Registry GUI” on page 400</a> |

The Oracle Solaris Product Registry is a tool to help you manage installed software. After you have installed the software, Product Registry provides a list of all the installed software by using the Oracle Solaris installation GUI or the `pkgadd` command.

You can use the Oracle Solaris Product Registry in a GUI or with a command-line interface (CLI). For more information on how to use the Oracle Solaris Product Registry CLI, see [“Managing Software With the Oracle Solaris Product Registry Command-Line Interface \(Task Map\)” on page 401](#).

The Oracle Solaris Product Registry GUI interface enables you to do the following:

- View a list of installed and registered software and some software attributes.
- View all Oracle Solaris system products that you installed in their localized version in the System Software Localizations directory.
- Find and launch an installer.
- Install additional software products.
- Uninstall software and individual software packages.

The Oracle Solaris Product Registry GUI main window consists of three panes of information:

- Installed, registered, and removed software
- Standard attributes of the currently selected software
- Attributes that are customized and attributes that are internal to the registered software

## ▼ How to View Installed or Uninstalled Software Information With the Oracle Solaris Product Registry GUI

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in \*System Administration Guide: Security Services\*](#).

### 2 Start the Oracle Solaris Product Registry tool.

```
prodreg &
```

The Oracle Solaris Product Registry main window is displayed.

**3 Click the turner control that is located to the left of the System Registry directory in the Registered Software box.**

The turner control changes directions, from pointing right, to pointing downward. You can expand or collapse any item in the registry, except an item that has a text file icon to its left.

The Software Installed in Registered Software box always contains the following components:

- The configuration software group that you chose when you installed the Oracle Solaris release. Software groups that can be displayed include Reduced Network Support, Core, End User System Support, Developer System Support, Entire Distribution, or Entire Distribution Plus OEM Support.
- Additional system software, which contains Oracle Solaris products that are not part of the software group you chose.
- Unclassified software that is not an Oracle Solaris product or part of the software group. This software includes any package that you installed by using the `pkgadd` command.

**4 Select directories until you find a software application to view.**

The list expands as you open directories.

**5 To view the attributes, select a directory or file.**

The Product Registry displays attribute information in the System Registry box.

- For software products that were installed with the installation GUI, the Oracle Solaris Product Registry contains values for at least the following: Title, Version, Location, and Installed on. Items in an expanded list under a product or software group inherit the version information of the product.
- If all or part of the product was removed with the `pkgrm` command, a cautionary icon appears next to the software product's name.

## ▼ How to Install Software With the Oracle Solaris Product Registry GUI

You can use Oracle Solaris Product Registry to find software and launch the installation GUI program. This program takes you through the installation of that software.

**1 Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

**2 Start the Oracle Solaris Product Registry tool.**

```
prodreg
```

The Oracle Solaris Product Registry main window is displayed.

- 3 **Depending on whether you are installing from a CD, a DVD, or from the network, do one of the following:**
  - **If you running a release that uses CDs, insert the CD into the CD-ROM drive.**
  - **If you are installing from a DVD, insert the DVD into the DVD-ROM drive.**
  - **If you are installing from the network, locate the net image of the software that you want to install.**

---

**Note** – In this Oracle Solaris release, the software is delivered on DVD *only*.

---

- 4 **To view the list of installed and registered software, click the turner control.**
- 5 **Click the New Install button at the bottom of the Oracle Solaris Product Registry window.**

The Select Installer dialog box is displayed. This box initially points to the /cdrom directory or the directory you are in.
- 6 **Select directories to find the Oracle Solaris installation GUI installer.**

Oracle Solaris installation GUI installers are located in various directories on the CDs and on the DVD.

  - Solaris 10 Software CDs or DVD.
  - Solaris 10 Documentation DVD.
  - Solaris 10 Languages CD. The installation GUI automatically starts when the CD is inserted.
- 7 **When you find the installer you want, select its name in the Files box.**
- 8 **Click OK.**

The installer you selected is launched.
- 9 **Follow the directions that are displayed by the installer to install the software.**

## ▼ **How to Uninstall Software With the Oracle Solaris Product Registry GUI**

- 1 **Become superuser or assume an equivalent role.**

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.



**2 Start the Oracle Solaris Product Registry tool.**

```
prodreg
```

The Oracle Solaris Product Registry main window is displayed.

**3 To view the list of installed and registered software, click the turner control.****4 Select directories until you find the name of the software that you want to uninstall.****5 Read the software attributes to make sure that this software is the software that you want to uninstall.****6 Click the Uninstall *software-product-name* button at the bottom of the Oracle Solaris Product Registry window.**

The software product you selected is uninstalled.

## Managing Software With the Oracle Solaris Product Registry Command-Line Interface (Task Map)

The following task map describes the software management tasks that you can perform with the Oracle Solaris Product Registry command-line interface.

| Task                                            | Description                                                                                                                               | For Instructions                                                                                  |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|
| View installed or uninstalled software.         | You can view software information by using the <code>browse</code> subcommand.                                                            | <a href="#">“How to View Installed or Uninstalled Software Information (prodreg)” on page 402</a> |
| View software attributes.                       | You can view specific software attributes by using the <code>info</code> subcommand.                                                      | <a href="#">“How to View Software Attributes (prodreg)” on page 405</a>                           |
| Check dependencies between software components. | You can view the components that depend on a specific software component by using the <code>info</code> subcommand.                       | <a href="#">“How to Check for Software Dependencies (prodreg)” on page 407</a>                    |
| Identify damaged software products.             | If you remove installed software files or packages without using the appropriate uninstaller, you can damage the software on your system. | <a href="#">“How to Identify Damaged Software Products (prodreg)” on page 408</a>                 |
| Uninstall software                              | You can remove software from your system by using the <code>uninstall</code> subcommand.                                                  | <a href="#">“How to Uninstall Software (prodreg)” on page 410</a>                                 |

| Task                                   | Description                                                                                                                                                                         | For Instructions                                                                     |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Uninstall damaged software.            | Uninstalling a damaged software component might fail if the uninstaller program for the software component has been removed from the system.                                        | <a href="#">“How to Uninstall Damaged Software (prodreg)”</a> on page 412            |
| Reinstall damaged software components. | If other software depends on a damaged software component, you might want to reinstall the damaged component, rather than uninstall the component and the other dependent software. | <a href="#">“How to Reinstall Damaged Software Components (prodreg)”</a> on page 414 |

## Managing Software With the Oracle Solaris Product Registry Command-Line Interface

The `prodreg` command is the command-line interface (CLI) to the Oracle Solaris Product Registry. The `prodreg` command supports several subcommands that enable you to manage the software on your system.

You can use the `prodreg` command in a terminal window to perform the following tasks:

- View a list of installed and registered software and software attributes.
- View all Oracle Solaris system products that you installed in their localized version in the System Software Localizations directory.
- Identify damaged software.
- Remove software entries from the Oracle Solaris Product Registry.
- Uninstall software and individual software packages.

For more information on how to manage the Oracle Solaris Product Registry by using the command-line interface, see the [prodreg\(1M\)](#) man page.

### ▼ How to View Installed or Uninstalled Software Information (prodreg)

You can view information about software in the Oracle Solaris Product Registry in a terminal window by using the `browse` subcommand to the `prodreg` command.

- 1 **Open a terminal window.**

## 2 Browse the Oracle Solaris Product Registry.

```
% prodreg browse
 BROWSE # +/-/. UUID # NAME
 ===== ===== =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
 3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software
```

The browse subcommand to the prodreg command displays the following information about registered software.

**BROWSE #** When you use the prodreg browse command, the Oracle Solaris Product Registry generates a *browse number* for each registered software component. This number can be used as an argument to either the prodreg browse command or the info subcommand to descend the hierarchy of specific registered components.

---

**Note** – Browse numbers might change when you reboot or reinstall your system. Do not store browse numbers in scripts or attempt to reuse them between separate login sessions.

---

**+/-/.** This field indicates if a software component has additional software component children registered in the Oracle Solaris Product Registry.

The following characters are displayed in this field:

- + indicates that the software component has additional children components that are not currently displayed.
- - indicates that the software component has additional children components that are currently displayed.
- . indicates that the software component does not have children components.

**UUID** This field lists the software's unique identifier in the Oracle Solaris Product Registry.

**#** This field indicates the *instance number* of the software component on the system. If the system contains multiple instances of a software component, the Oracle Solaris Product Registry assigns a separate instance number to each instance of the component.

**NAME** This field lists the localized name of the software. The name of the Oracle Solaris OS in this sample output is the Oracle Solaris 10 system software.

**3 Browse the information for one of the software components that are listed in the Oracle Solaris Product Registry.**

```
% prodreg browse -m "name"
```

The -m "*name*" command displays information on the software component with the name *name*.

**4 If the system contains multiple instances of *name* software, type the following command to browse the Oracle Solaris Product Registry:**

```
% prodreg browse -u name-UUID -i
instance -n number
```

```
-u name-UUID
```

Displays information on the *name* software component with the unique identifier *name-UUID*.

```
-i instance
```

Displays information on the *name* software component with the instance number *instance*.

```
-n number
```

Displays software information by referencing the component's browse number *number*.

**5 Repeat Step 3 and Step 4 for each software component that you want to browse.**

**Example 21-1 Viewing Software Information by Component Name (prodreg)**

The following example shows how to view software information by referencing the component's name.

```
% prodreg browse
 BROWSE # +/-/. UUID # NAME
 =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
 3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

% prodreg browse -m "Solaris 10 System Software"
```

**Example 21-2 Viewing Software Information by Component Browse Number (prodreg)**

The following example shows how to use the -n option with the prodreg browse command to view software information by referencing the component's browse number.

```
% prodreg browse
 BROWSE # +/-/. UUID # NAME
```

```

===== ===== ===== ===== =====
1 - root 1 System
 Registry
2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

```

```
% prodreg browse -n 2
```

### Example 21-3 Viewing Software Information by Component UUID (prodreg)

The following example shows how to use the `-u` option with the `prodreg browse` command to view software information by referencing the component's UUID. The UUID is the software's unique identifier in the Oracle Solaris Product Registry.

```

% prodreg browse
 BROWSE # +/-/. UUID # NAME
===== ===== ===== ===== =====
1 - root 1 System
 Registry
2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

```

```
% prodreg browse -u a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b
```

## ▼ How to View Software Attributes (prodreg)

You can view specific software attributes by using the `info` subcommand of the `prodreg` command.

The `prodreg info` command displays a variety of information about registered software, including the following items:

- Software component name
- Software component description
- Required components of the software
- Other components that require the software
- Base directory of the software
- Path to the software component

### 1 Open a terminal window.

## 2 Browse the Oracle Solaris Product Registry.

```
% prodreg browse
 BROWSE # +/-/. UUID # NAME
 =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
 3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software
```

## 3 View the attributes for one of the listed software components.

```
% prodreg info -m "name"
```

The -m “name” command displays the attributes of the software component with the name *name*.

## 4 Repeat [Step 3](#) for each software component you want to view.

### Example 21–4 Viewing Software Attributes by Component Name (prodreg)

The following example shows how to view software attributes by referencing the component's name.

```
% prodreg browse
 BROWSE # +/-/. UUID # NAME
 =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
 3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

% prodreg info -m "Solaris 10 System Software"
```

### Example 21–5 Viewing Software Attributes by Component Browse Number (prodreg)

The following example shows how to use the -n option with the prodreg info command to view software attributes by referencing the component's browse number.

```
% prodreg browse
 BROWSE # +/-/. UUID # NAME
 =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
```

```

3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

% prodreg info -n 2

```

### Example 21-6 Viewing Software Attributes by Component UUID (prodreg)

The following example shows how to use the `-u` option with the `prodreg info` command to view software attributes by referencing the component's UUID. The UUID is the software's unique identifier in the Oracle Solaris Product Registry.

```

% prodreg browse
 BROWSE # +/-/. UUID # NAME
 ===== ===== =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
 3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

% prodreg info -u a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b

```

## ▼ How to Check for Software Dependencies (prodreg)

You can use the `prodreg info` command to view components that depend on a specific software component. You might want to check dependencies between software products before you uninstall specific components.

- 1 Open a terminal window.
- 2 Browse the Oracle Solaris Product Registry.

```

% prodreg browse
 BROWSE # +/-/. UUID # NAME
 ===== ===== =====
 1 - root 1 System
 Registry
 2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
 System
 Software
 3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
 Software

```

Repeat the `prodreg browse` command until the software component you want to check is displayed. See [“How to View Installed or Uninstalled Software Information \(prodreg\)” on page 402](#) for more information on browsing the Oracle Solaris Product Registry by using the `prodreg browse` command.

### 3 View the dependencies of a specific software component.

```
% prodreg info -m "name" -a "Dependent Components"
```

```
-m "name"
```

Displays the attributes of the software component with the name *name*.

```
-a "Dependent Components"
```

Displays components that depend on *name* software by displaying the values of the Dependent Components attribute.

This command output lists the software components that depend on *name* software.

#### Example 21-7 Viewing Components That Depend on Other Software Products (prodreg)

The following example shows how to view the components that depend on the software product that is named ExampleSoft.

```
% prodreg -m "ExampleSoft" -a "Dependent Components"
```

```
Dependent Components:
```

| Name         | UUID                                 | # |
|--------------|--------------------------------------|---|
| -----        | -----                                | - |
| ExampleSoftA | 7f49ecvb-1ii2-11b2-a3f1-0800119u7e8e | 1 |

## ▼ How to Identify Damaged Software Products (prodreg)

If you remove installed software files or packages without using the appropriate uninstaller, you can damage the software on your system. If software is damaged, the software might not function properly. You can use the `info` subcommand of the `prodreg` command to help you determine if a software product is damaged.

### 1 View the Oracle Solaris Product Registry information on the software you want to check.

```
% prodreg browse -m name
```

### 2 Verify that the software component is damaged.

```
% prodreg info -u name-UUID -i 1 -d
```

### 3 Identify the packages that form the *name-UUID* software component.

```
% prodreg info -u name-UUID -i 1 -a PKGS
```

The output of this command might list more than one package.



- 4 Verify that the packages displayed in the previous step are installed on the system by running the `pkginfo` command for each package.

```
% pkginfo component-a-pkg
% pkginfo component-b-pkg
.
.
.
```

### Example 21-8 Identifying Damaged Software Components (prodreg)

The following example shows how to determine if the ExampleSoft software component is damaged.

```
% prodreg browse -m Examplesoft
BROWSE # +/-/. UUID # NAME
=====
1 - root 1 System
Registry
2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
System
Software
3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
Software
4 - 95842091-725a-8501-ef29-0472985982be 1 ExampleSoft
233 . 90209809-9785-b89e-c821-0472985982be 1 Example Doc
234 . EXS0zzt 1
235 . EXS0blob 1 Example Data
```

The ExampleSoft child component EXS0zzt does not have an entry in the NAME field. The ExampleSoft software might be damaged. You would use the `prodreg info` command with the `-u`, `-i`, and `-d` options to determine if the ExampleSoft software is damaged.

```
% prodreg info -u 95842091-725a-8501-ef29-0472985982be -i 1 -d
isDamaged=TRUE
```

The output `isDamaged=TRUE` indicates that the ExampleSoft software is damaged. You would use the `-a PKGS` option of the `prodreg info` command to identify the ExampleSoft software packages.

```
% prodreg info
-u 95842091-725a-8501-ef29-0472985982be
-i 1 -a PKGS
pkgs:
EXS0zzt EXS0blob
```

To verify that the EXS0zzt and EXS0blob packages are installed on the system, you would use the `pkginfo` command.

```
% pkginfo EXS0zzt
ERROR: information for "EXS0zzt" was not found
```

```
% pkginfo EXS0blob
application EXS0blob Example Data
```

The output of the `pkginfo` command indicates that the `EXS0zzt` package is not installed on the system. Thus, the `ExampleSoft` software is damaged.

## ▼ How to Uninstall Software (prodreg)

You can use the `uninstall` subcommand of the `prodreg` command to remove software from your system. When you uninstall software by using the `prodreg uninstall` command, you remove a specified software and all the child components associated with that software. Before you remove software, verify that other software does not depend on the software you want to uninstall. See [“How to Check for Software Dependencies \(prodreg\)” on page 407](#).

After you uninstall software, you can remove that software and all the child components of that software from the Oracle Solaris Product Registry by using the `prodreg unregister -r` command.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in \*System Administration Guide: Security Services\*](#).

### 2 View the information on the software you want to uninstall.

```
prodreg browse -u name-UUID
```

### 3 Uninstall the software.

```
prodreg uninstall -u name-UUID
```

### 4 Check the dependencies for the software that you want to uninstall.

```
prodreg info -u name-UUID
```

Check the following information in the output of the `prodreg info` command.

- Child Components – Lists the software components that are associated with the *name* software component. When you unregister the *name* software, you also unregister the child components of *name* software. If the output of the previous `prodreg info` command lists any child components, verify that you want to unregister these child components.
- Required Components – Lists the software components that are required by the *name* software component. Software components might require other components that are not child components. When you uninstall and unregister a component, only child components are unregistered and uninstalled.

- **Dependent Components** – Lists the components that require *name* software to run. When you unregister the *name* software, you also unregister the dependent components of *name* software. If the output of the `prodreg info` command lists any dependent components, verify that you want to unregister these dependent components.

In the previous sample output, *name* software does not have any dependent components.

##### 5 Check the dependencies of *name* software's child components.

```
prodreg info -u component-a-UUID -i 1 -a "Dependent Components"
```

The sample output shows that no other software depends on the child components of *name* software.

##### 6 Unregister the software and its child components.

```
prodreg unregister -r -u name-UUID -i 1
```

- r** Recursively unregisters software with the unique identifier *name-UUID* and all the child components of this software.
- u name-UUID** Specifies the unique identifier of the software you want to unregister.
- i 1** Specifies the instance of the software you want to unregister.

### Example 21–9 Uninstalling Software Components (prodreg)

The following example shows how to uninstall ExampleSoft software and all the child components of ExampleSoft software.

```
prodreg browse -m "ExampleSoft"
BROWSE # +/-/. UUID # NAME
=====
1 - root 1 System
Registry
2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
System
Software
3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
Software
1423 - 95842091-725a-8501-ef29-0472985982be 1 ExampleSoft
1436 . 90209809-9785-b89e-c821-0472985982be 1 Example Doc
1437 - EXS0zzt 1 Example Data
1462 . EXS0blob 1 Example Data

prodreg uninstall -u 95842091-725a-8501-ef29-0472985982be -i 1

prodreg info -u 95842091-725a-8501-ef29-0472985982be
Title: ExampleSoft Software
.
```

```

.
.
Child Components:
Name UUID #

Example Doc 90209809-9785-b89e-c821-0472985982be 1
Example Data EXS0zzt 1

Required Components:
Name UUID #

Example Doc 90209809-9785-b89e-c821-0472985982be 1
Example Data EXS0zzt 1

prodreg info -u 90209809-9785-b89e-c821-0472985982be -i 1
 -a "Dependent Components"
Dependent Components:
Name UUID #

ExampleSoft 95842091-725a-8501-ef29-0472985982be 1

prodreg info -u EXS0zzt -i 1 -a "Dependent Components"
Dependent Components:
Name UUID #

ExampleSoft 95842091-725a-8501-ef29-0472985982be 1

prodreg info -u EXS0blob -i 1 -a "Dependent Components"
Dependent Components:
Name UUID #

Example Data EXS0zzt 1

prodreg unregister -r -u 95842091-725a-8501-ef29-0472985982be -i 1

```

## ▼ How to Uninstall Damaged Software (prodreg)

If you try to uninstall a damaged software component by using the `prodreg uninstall` command, the command might fail. This failure can occur if the uninstaller program for the software component has been removed from the system.

Follow these steps to uninstall a software component with no associated uninstaller program on the system.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.

### 2 View the information on the software you want to uninstall.

```
prodreg browse -m "name"
```

### 3 Uninstall the software.

```
prodreg uninstall -u UUID -i 1
```

### 4 Identify the uninstaller program for the software component.

```
prodreg info -m "name" -a uninstallprogram
```

### 5 Determine if the uninstaller is in the registered location.

```
ls uninstaller-location uninstaller-location
```

### 6 Remove the software from the system in one of the following ways:

- If you have a system backup available, follow these steps:

- a. Load the uninstaller program from the backup.

- b. Run the uninstaller program from a shell command-line interface such as a terminal window.

- If you do not have access to the uninstaller program on a backup, follow these steps:

- a. Unregister the software component.

```
prodreg unregister -u UUID -i 1
```

- b. Remove any remaining registered components that are required by the software you want to remove.

```
pkgrm component-a-UUID
```

## Example 21–10 Uninstalling Damaged Software (prodreg)

The following example shows how to uninstall the damaged ExampleSoft software. In this example, the uninstaller program is not readily available on a system backup.

```
prodreg browse -m Examplesoft
BROWSE # +/-/. UUID # NAME
=====
1 - root 1 System
Registry
2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
System
Software
3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
Software
4 - 95842091-725a-8501-ef29-0472985982be 1 ExampleSoft
233 . 90209809-9785-b89e-c821-0472985982be 1 Example Doc
234 . EXS0zzt 1
235 . EXS0blob 1 Example Data
```

```
prodreg uninstall -u 95842091-725a-8501-ef29-0472985982be -i 1
The install program requested could not be found

prodreg info -m "ExampleSoft" -a uninstallprogram
uninstallprogram: /usr/bin/java -mx64m -classpath
/var/sadm/prod/org.example.ExampleSoft/987573587 uninstall_ExampleSoft

ls /var/sadm/prod/org.example.ExampleSoft/987573587
/var/sadm/prod/org.example.ExampleSoft/987573587:
No such file or directory

prodreg unregister -u 95842091-725a-8501-ef29-0472985982be -i 1

pkgrm EXS0blob
```

## ▼ How to Reinstall Damaged Software Components (prodreg)

If other software depends on a damaged software component, you might want to reinstall the damaged component, rather than uninstall the component and the other dependent software. You can use the `-f` option with the `prodreg unregister` command to forcibly the unregister the damaged component. Then, you can reinstall the component.

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)”](#) in *System Administration Guide: Security Services*.

### 2 View the information on the software you want to reinstall.

```
prodreg browse -m "name"
```

### 3 Identify the software that depends on the software you want to reinstall.

```
prodreg info -m "name" -a "Dependent Components"
```

### 4 Unregister the software component you want to reinstall.

```
prodreg unregister -f -u UUID
```

### 5 Reinstall the software component.

```
/usr/bin/java -cp /usr/installers/installer
```

The `installer` option specifies the name of the installer program for `name` software.

## Example 21–11 Reinstalling Damaged Software Components (prodreg)

The following example shows how to reinstall the damaged software component `ComponentSoft` without unregistering or uninstalling the dependent component `ExampleSoft`.

```

prodreg browse -m "ComponentSoft"
BROWSE # +/-/. UUID # NAME
=====
1 - root 1 System
Registry
2 + a01ee8dd-1dd1-11b2-a3f2-0800209a5b6b 1 Solaris 10
System
Software
3 + 8f64eabf-1dd2-11b2-a3f1-0800209a5b6b 1 Unclassified
Software
4 . 86758449-554a-6531-fe90-4352678362fe 1 ComponentSoft

prodreg info -m "ComponentSoft" -a "Dependent Components"
Dependent Components:
Name UUID #

ExampleSoft 95842091-725a-8501-ef29-0472985982be 1

prodreg unregister -f -u 86758449-554a-6531-fe90-4352678362fe -i 1

/usr/bin/java -cp /usr/installers/org.example.componentsoft

```





# Managing Software by Using Oracle Solaris Package Commands (Tasks)

---

This chapter describes how to add, verify, and remove software packages by using the Oracle Solaris package commands. For information on the procedures associated with performing these tasks, see [“Managing Software Packages by Using Package Commands \(Task Map\)”](#) on page 417.

## Managing Software Packages by Using Package Commands (Task Map)

The following task map describes the software management tasks that you can perform with the package commands.

| Task                                                    | Description                                                                                           | For Instructions                                                                                 |
|---------------------------------------------------------|-------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Add software packages to the local system.              | You can add software packages to the local system by using the <code>pkgadd</code> command.           | <a href="#">“How to Add Software Packages (pkgadd)”</a> on page 418                              |
| Add software packages to a spool directory.             | You can add software packages to a spool directory without actually installing the software.          | <a href="#">“Adding a Software Package to a Spool Directory”</a> on page 421                     |
| List information about all installed software packages. | You can list information about installed packages by using the <code>pkginfo</code> command.          | <a href="#">“How to List Information About All Installed Packages (pkginfo)”</a> on page 423     |
| Check the integrity of installed software packages.     | You can verify the integrity of installed software packages by using the <code>pkgchk</code> command. | <a href="#">“How to Check the Integrity of Installed Software Packages (pkgchk)”</a> on page 424 |

| Task                                        | Description                                                                                                                                                                                                                                                              | For Instructions                                                                                                                 |
|---------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------|
| Check the integrity of an installed object. | You can verify the integrity of an installed object by using the <code>pkchk</code> command with the <code>-p</code> and <code>-P</code> options. The <code>-p</code> option specifies the full path name. The new <code>-P</code> option specifies a partial path name. | “ <a href="#">How to Check the Integrity of Installed Objects ( <code>pkchk -p</code>, <code>pkchk -P</code> )</a> ” on page 425 |
| Remove software packages.                   | You can remove unneeded software packages by using the <code>pkgrm</code> command.                                                                                                                                                                                       | “ <a href="#">How to Remove Software Packages ( <code>pkgrm</code> )</a> ” on page 427                                           |

## Using Package Commands to Manage Software Packages

The following procedures explain how to manage software packages by using package commands.

### ▼ How to Add Software Packages (`pkgadd`)

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

#### 2 Remove any already installed packages with the same names as the packages you are adding.

This step ensures that the system keeps a proper record of software that has been added and removed. Sometimes, you might want to maintain multiple versions of the same application on the system. For strategies on maintaining multiple software copies, see “[Guidelines for Removing Packages \( `pkgrm` \)](#)” on page 392. For task information, see “[How to Remove Software Packages \( `pkgrm` \)](#)” on page 427.

#### 3 Add a software package to the system.

```
pkgadd -a admin-file
-d device-name pkgid ...
```

`-a admin-file` (Optional) Specifies an administration file that the `pkgadd` command should check during the installation. For details about using an administration file, see “[Using an Administration File](#)” on page 393.

`-d device-name` Specifies the absolute path to the software packages. `device-name` can be the path to a device, a directory, or a spool directory. If you do not specify the path where the package resides, the `pkgadd` command checks the default spool directory (`/var/spool/pkg`). If the package is not there, the package installation fails.

*pkgid* (Optional) Is the name of one or more packages, separated by spaces, to be installed. If omitted, the `pkgadd` command installs all available packages from the specified device, directory, or spool directory.

If the `pkgadd` command encounters a problem during installation of the package, it displays a message related to the problem, followed by this prompt:

```
Do you want to continue with this installation?
```

Respond with `yes`, `no`, or `quit`. If more than one package has been specified, type `no` to stop the installation of the package being installed. The `pkgadd` command continues to install the other packages. Type `quit` to stop the installation.

#### 4 Verify that the package has been installed successfully.

```
pkgchk -v pkgid
```

If no errors occur, a list of installed files is returned. Otherwise, the `pkgchk` command reports the error.

### Example 22-1 Adding Software Packages From a Mounted CD

The following example shows how install the `SUNWpl5u` package from a mounted Oracle Solaris 10 CD. The example also shows how to verify that the package files were installed properly.

```
pkgadd -d /cdrom/cdrom0/Solaris_10/Product SUNWpl5u
.
.
.
Installation of <SUNWpl5u> was successful.
pkgchk -v SUNWpl5u
/usr
/usr/bin
/usr/bin/perl
/usr/perl5
/usr/perl5/5.8.4
.
.
.
```

This example shows the path to use if you are *not* running at least the Solaris 10 10/08 release.

```
pkgadd -d /cdrom/cdrom0/s0/Solaris_10/Product SUNWpl5u
.
.
.
Installation of <SUNWpl5u> was successful.
pkgchk -v SUNWpl5u
/usr
/usr/bin
```

```
/usr/bin/perl
/usr/perl5
/usr/perl5/5.8.4
.
.
.
```

### Example 22-2 Installing Software Packages From a Remote Package Server

If the packages you want to install are available from a remote system, you can manually mount the directory that contains the packages (in package format) and install packages on the local system.

The following example shows how to install software packages from a remote system. In this example, assume that the remote system named `package-server` has software packages in the `/latest-packages` directory. The `mount` command mounts the packages locally on `/mnt`. The `pkgadd` command installs the `SUNWpl5u` package.

```
mount -F nfs -o ro package-server:/latest-packages /mnt
pkgadd -d /mnt SUNWpl5u
```

```
.
.
.
Installation of <SUNWpl5u> was successful.
```

If the automounter is running at your site, you do not need to mount the remote package server manually. Instead, use the automounter path, in this case, `/net/package-server/latest-packages`, as the argument to the `-d` option.

```
pkgadd -d /net/package-server/latest-packages SUNWpl5u
```

```
.
.
.
Installation of <SUNWpl5u> was successful.
```

### Example 22-3 Installing Software Packages From a Remote Package Server by Specifying an Administration File

This example is similar to the previous example, except that it uses the `-a` option and specifies an administration file named `noask-pkgadd`, which is illustrated in [“Avoiding User Interaction When Adding Packages \(pkgadd\)” on page 393](#). In this example, assume that the `noask-pkgadd` administration file is in the default location, `/var/sadm/install/admin`.

```
pkgadd -a noask-pkgadd -d /net/package-server/latest-packages SUNWpl5u
```

```
.
.
.
Installation of <SUNWpl5u> was successful.
```

**Example 22–4** Installing Software Packages From an HTTP URL

The following example shows how to install a package using an HTTP URL as the device name. The URL must point to a stream-formatted package.

```
pkgadd -d http://install/xf86-4.3.0-video.pkg
Downloading...
.....25%.....50%.....75%.....100%
Download Complete
```

The following packages are available:

```
1 SUNWxf86r XFree86 Driver Porting Kit (Root)
 (i386) 4.3.0,REV=0.2003.02.28
2 SUNWxf86u XFree86 Driver Porting Kit (User)
 (i386) 4.3.0,REV=0.2003.02.28
```

```
.
```

```
.
```

```
.
```

## Adding a Software Package to a Spool Directory

For convenience, you can copy frequently installed packages to a spool directory. If you copy packages to the default spool directory, `/var/spool/pkg`, you do not need to specify the source location of the package (`-d device-name` argument) when you use the `pkgadd` command. The `pkgadd` command, by default, checks the `/var/spool/pkg` directory for any packages that are specified on the command line. Note that copying packages to a spool directory is not the same as installing the packages on a system.

### ▼ How to Add Software Packages to a Spool Directory (`pkgadd`)

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

#### 2 Remove any already spooled packages with the same names as the packages you are adding.

For information on removing spooled packages, see [Example 22–16](#).

#### 3 Add a software package to a spool directory.

```
pkgadd -d device-name
-s spooldir pkgid ...
```

`-d device-name` Specifies the absolute path to the software packages. `device-name` can be the path to a device, a directory, or a spool directory.

- `-s spooldir` Specifies the name of the spool directory where the package will be spooled. You must specify a *spooldir*.
- `pkgid` (Optional) Is the name of one or more packages, separated by spaces, to be added to the spool directory. If omitted, the `pkgadd` command copies all available packages.

#### 4 Verify that the package has been copied successfully to the spool directory.

```
$ pkginfo -d spooldir
| grep pkgid
```

If *pkgid* was copied correctly, the `pkginfo` command returns a line of information about the *pkgid*. Otherwise, the `pkginfo` command returns the system prompt.

### Example 22-5 Setting Up a Spool Directory From a Mounted CD

The following example shows how to transfer the SUNWman package from a mounted SPARC based Oracle Solaris 10 CD to the default spool directory (`/var/spool/pkg`).

```
pkgadd -d /cdrom/cdrom0/Solaris_10/Product -s /var/spool/pkg SUNWman
```

Transferring <SUNWman> package instance

### Example 22-6 Setting Up a Spool Directory From a Remote Software Package Server

If packages you want to copy are available from a remote system, you can manually mount the directory that contains the packages, in package format, and copy them to a local spool directory.

The following example shows the commands for this scenario. In this example, assume that the remote system named `package-server` has software packages in the `/latest-packages` directory. The `mount` command mounts the package directory locally on `/mnt`. The `pkgadd` command copies the `SUNWp15p` package from `/mnt` to the default spool directory (`/var/spool/pkg`).

```
mount -F nfs -o ro package-server:/latest-packages /mnt
pkgadd -d /mnt -s /var/spool/pkg SUNWp15p
Transferring <SUNWp15p> package instance
```

If the automounter is running at your site, you do not have to mount the remote package server manually. Instead, use the automounter path, in this case, `/net/package-server/latest-packages`, as the argument to the `-d` option.

```
pkgadd -d /net/package-server/latest-packages -s /var/spool/pkg SUNWp15p
```

Transferring <SUNWp15p> package instance

**Example 22-7** Installing Software Packages From the Default Spool Directory

The following example shows how to install the SUNWpl5p package from the default spool directory. When no options are used, the `pkgadd` command searches the `/var/spool/pkg` directory for the named packages.

```
pkgadd SUNWpl5p
.
.
.
Installation of <SUNWpl5p> was successful.
```

## ▼ How to List Information About All Installed Packages (pkginfo)

- List information about installed packages by using the `pkginfo` command.

```
$ pkginfo
```

**Example 22-8** Listing Installed Packages

This example shows how to list all packages installed on a local system, whether that system is a stand-alone system or a server. The output shows the primary category, package name, and the description of the package.

```
$ pkginfo
system SUNWaccr System Accounting, (Root)
system SUNWaccu System Accounting, (Usr)
system SUNWadmap System administration applications
system SUNWadmc System administration core libraries
.
.
.
```

**Example 22-9** Displaying Detailed Information About Software Packages

This example shows how to list all packages installed on a system by specifying the long format, which includes all available information about the designated packages.

```
$ pkginfo -l SUNWcar
PKGINST: SUNWcar
NAME: Core Architecture, (Root)
CATEGORY: system
ARCH: i386.i86pc
VERSION: 11.10.0,REV=2005.01.21.16.34
BASEDIR: /
VENDOR: Oracle Corporation
DESC: core software for a specific hardware platform group
```

```
PSTAMP: on10ptchfeatx20110211045100
INSTDATE: Mar 03 2011 10:57
HOTLINE: Please contact your local service provider
STATUS: completely installed
FILES: 2 installed pathnames
 2 shared pathnames
 2 directories
```

## ▼ How to Check the Integrity of Installed Software Packages (pkgchk)

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see [“Configuring RBAC \(Task Map\)” in \*System Administration Guide: Security Services\*](#).

### 2 Check the status of an installed package.

- To check the file attributes and contents, type the following:

```
pkgchk -a| -c -v pkgid ...
```

- To specify the absolute path of the spool directory, type the following:

```
pkgchk -d spooldir pkgid ...
```

- a Specifies to audit only the file attributes (the permissions), rather than the file attributes and the contents, which is the default.
- c Specifies to audit only the file contents, rather than the file contents and attributes, which is the default.
- v Specifies verbose mode, which displays file names as they are processed.
- d *spooldir* Specifies the absolute path of the spool directory.
- pkgid* (Optional) Is the name of one or more packages, separated by spaces. If you do not specify a *pkgid*, all the software packages installed on the system are checked.

### Example 22–10 Checking the Contents of Installed Software Packages

The following example shows how to check the contents of a package.

```
pkgchk -c SUNWbash
```

If no errors occur, the system prompt is returned. Otherwise, the pkgchk command reports the error.



**Example 22–11** Checking the File Attributes of Installed Software Packages

The following example shows how to check the file attributes of a package.

```
pkgchk -a SUNWbash
```

If no errors occur, the system prompt is returned. Otherwise, the `pkgchk` command reports the error.

**Example 22–12** Checking Software Packages Installed in a Spool Directory

The following example shows how to check a software package that was copied to a spool directory (`/export/install/packages`).

```
pkgchk -d /export/install/packages
checking spooled package <SUNWadmap>
checking spooled package <SUNWadmfw>
checking spooled package <SUNWadmC>
checking spooled package <SUNWsadml>
```

The checks made on a spooled package are limited because not all information can be audited until a package is installed.

## ▼ How to Check the Integrity of Installed Objects (`pkgchk -p`, `pkgchk -P`)

This procedure explains how to use the `pkgchk` command to check the integrity of installed objects. The new `-P` option enables you to specify a partial path. This option has been added to assist you in mapping files to packages. Use this option with the `-l` option to list the information about the files that contain the partial path. Use the `-p` option to check the integrity of installed objects by specifying the full path. For more information, see the [`pkgchk\(1M\)` man page](#).

### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “[Configuring RBAC \(Task Map\)](#)” in *System Administration Guide: Security Services*.

### 2 Check the integrity of an installed object.

- To verify the integrity of an installed object for a full path name or path names, type the following:

```
pkgchk -lp path-name
```

- To verify the integrity of an installed object for a partial-path name or path names, type the following:

- `# pkgchk -lP partial-path-name`
- `-p path` Checks the accuracy only of the path name or path names that are listed. Path can be one or more path names separated by commas. Specifies to audit only the file attributes (the permissions), rather than the file attributes and the contents, which is the default.
- `-P partial-path` Checks the accuracy of only the partial path name or path names that are listed. The partial-path can be one or more partial path names separated by commas. Matches any path name that contains the string contained in the partial path. Specifies to audit only the file contents, rather than the file contents and attributes, which is the default.
- `-l` Lists information about the selected files that make up a package. This option is not compatible with the `-a`, `-c`, `-f`, `-g`, and `-v` options. Specifies verbose mode, which displays file names as they are processed.

### Example 22-13 Checking the Integrity of an Installed Object by Specifying a Full Path Name

This example shows you how to use the `pkgchk -lp` command to check the contents/attributes of an object on a file system by a specifying the full path name. The `-l` option lists information on the selected files that make up a package.

```
pkgchk -lp /usr/sbin/pkgadd
Pathname: /usr/sbin/pkgadd
Type: regular file
Expected mode: 0555
Expected owner: root
Expected group: sys
Expected file size (bytes): 867152
Expected sum(1) of contents: 45580
Expected last modification: Jul 02 02:20:34 2004
Referenced by the following packages:
 SUNWpkgcmdsu
Current status: installed
```

### Example 22-14 Checking the Integrity of an Installed Object by Specifying a Partial Path Name

This example shows you how to use the `pkgchk -lP` command to check the contents/attributes of an object on a file system by a specifying a partial path name, such as a file or directory name. The `-l` option lists information on the selected files that make up a package.

```
pkgchk -lP /sbin/pkgadd
Pathname: /usr/sbin/pkgadd
Type: regular file
Expected mode: 0555
Expected owner: root
Expected group: sys
Expected file size (bytes): 867152
Expected sum(1) of contents: 45580
```

```

Expected last modification: Jul 02 02:20:34 2004
Referenced by the following packages:
 SUNWpkgcmdsu
Current status: installed

Pathname: /usr/sbin/pkgask
Type: linked file
Source of link: ../../usr/sbin/pkgadd
Referenced by the following packages:
 SUNWpkgcmdsu
Current status: installed

```

## Removing Software Packages

To remove or uninstall a software package, use the associated tool that you used to add or install a software package. For example, if you used the Oracle Solaris installation GUI to install software, use the Oracle Solaris installation GUI to uninstall software.




---

**Caution** – Do not use the `rm` command to remove software packages. Doing so will result in inaccuracies in the database that keeps track of all installed packages on the system.

---

### ▼ How to Remove Software Packages (`pkgrm`)

#### 1 Become superuser or assume an equivalent role.

Roles contain authorizations and privileged commands. For more information about roles, see “Configuring RBAC (Task Map)” in *System Administration Guide: Security Services*.

#### 2 Remove an installed package.

```
pkgrm pkgid ...
```

`pkgid` identifies the name of one or more packages, separated by spaces, to be removed. If omitted, the `pkgrm` command removes all available packages.

#### Example 22–15 Removing Software Packages

This example shows how to remove a package.

```
pkgrm SUNWctu
```

```

The following package is currently installed:
 SUNWctu Netra ct usr/platform links (64-bit)
 (sparc.sun4u) 11.9.0,REV=2001.07.24.15.53

```

```
Do you want to remove this package? y
```

```
Removing installed package instance <SUNWctu>
Verifying package dependencies.
Processing package information.
Removing pathnames in class <none>
.
.
.
```

### Example 22-16 Removing a Spooled Software Package

This example shows how to remove a spooled package.

```
pkgrm -s /export/pkg SUNWaudh
The following package is currently spooled:
 SUNWaudh Audio Header Files
 (sparc) 11.10.0,REV=2003.08.08.00.03
Do you want to remove this package? y
Removing spooled package instance <SUNWaudh>
```

# Managing Patches

---

Patch management involves *applying* patches and software updates to a system. Patch management might also involve removing unwanted or faulty patches. Removing patches is also called *backing out* patches.

This chapter provides the following information:

- “About Patches” on page 429
- “Patching Strategy” on page 430
- “Downloading a Patch” on page 432
- “Displaying Information About Patches” on page 433
- “Applying a Patch” on page 433
- “Removing a Patch” on page 434
- “Patch Management Terms and Definitions” on page 435

## About Patches

A patch is an accumulation of fixes for a known or potential problem within the Oracle Solaris operating system or other supported software. A patch can also provide a new feature or an enhancement to a particular software release. A patch consists of files and directories that replace or update existing files and directories. Therefore, patches are used for the following purposes:

- Delivering bug fixes
- Delivering new functionality
- Delivering new hardware support
- Delivering performance enhancements or enhancements to existing utilities

Patches are identified by unique patch IDs. A patch ID is an alphanumeric string that is a patch base code and a number that represents the patch revision number joined with a hyphen. For example, patch 119254-78 is the patch ID for the SunOS 5.10 kernel update patch, 78th revision.

## Patching Strategy

Depending on your requirements and the time available for the maintenance window, you can use any of the following patching strategies:

- “Live Upgrade” on page 430
- “Applying an Oracle Solaris Update or an Oracle Solaris Update Patch Bundle” on page 431
- “Applying a Recommended Patch Cluster” on page 431
- “Applying a Critical Patch Update” on page 432
- “Applying an Enterprise Installation Standards Patch Baseline” on page 432

---

**Note** – Firmware updates are not considered patches and they cannot be applied using the `patchadd` command. To apply a firmware update, follow the instructions in the README file of the corresponding firmware.

---

Before performing any patch related operations on an Oracle Solaris system, make sure that you have applied all of the Oracle recommended patching utility patches. You can download patches for the patching utilities, patch bundles, and patch clusters from the [My Oracle Support](#) (MOS) web site.

---

**Note** –

- The term “obsolete” used in the Sun SVR4 Patch Architecture equates to the term “superseded” in MOS.
  - The term “withdrawn” used in the Sun SVR4 Patch Architecture equates to the term “obsolete” in MOS.
  - In MOS, if a patch is marked “obsolete,” the patch has been withdrawn from the release.
  - In a patch, the `SUNW_OBSOLETES` field present in the `pkginfo` file specifies the patches superseded by the patch.
- 

## Live Upgrade

Oracle Solaris Live Upgrade (LU) provides a method of upgrading a system while the system continues to operate. While your current boot environment is running, you can duplicate the boot environment and then upgrade the duplicate. The original system configuration remains fully functional and is not affected by the upgrade or installation of an archive. When you are ready, you can activate the new boot environment by rebooting the system. If a failure occurs, you can quickly revert to the original boot environment with a simple reboot. This switch eliminates the normal downtime of the test and evaluation process. For more information, refer to *Oracle Solaris 10 8/11 Installation Guide: Solaris Live Upgrade and Upgrade Planning*.

## Applying an Oracle Solaris Update or an Oracle Solaris Update Patch Bundle

You can either install or upgrade to a new Oracle Solaris 10 update. An Oracle Solaris 10 Update is a full release image containing new features with all available patches pre-applied. Alternately you can apply the corresponding Oracle Solaris Update Patch Bundle. These patch bundles contain the equivalent set of patches that are included in the corresponding Oracle Solaris release image.

For example, consider the Oracle Solaris 10 5/08 Patch Bundle, and the corresponding release image Oracle Solaris 10 5/08. The Oracle Solaris 10 5/08 Patch Bundle contains the equivalent set of patches to the Oracle Solaris 10 5/08 release. The patch bundle does not include the new packages contained in the Oracle Solaris 10 5/08 release. Therefore, new features in the release that depend on new packages introduced in a release are not available in the patch bundle. However, any change to the pre-existing code is delivered in the patch bundle.

## Applying a Recommended Patch Cluster

The Solaris OS Recommended Patch Cluster contains all the available Oracle Solaris OS patches. These patches include:

- Fixes related to security
- Fixes related to data corruption
- Fixes related to system availability issues
- Recommended patches
- Latest patch utility patches
- Any other required patches

The recommended patch cluster bundle comes with install scripts (wrappers around the `patchadd` command) that perform the following functions:

- Filter out false negatives from the patch utility return codes. Only those errors that require further investigation by the user are reported.
- Exit as soon as an unexpected failure is encountered. This exit prevents problems that might occur by applying further patches.
- Include context intelligence for patching operations. The scripts inform the user when zones need to be halted, and provide phased installation to handle patches that absolutely require an immediate reboot before further patches can be applied.
- Provide better integration with Oracle Solaris Live Upgrade.
- Perform space checking prior to installing each patch. If there is insufficient space, installation is halted.

## Applying a Critical Patch Update

The Oracle Solaris OS critical patch update (CPU) is an archived snapshot of the Oracle Solaris OS recommended patch cluster. The standard practice of Oracle is to release a CPU once every quarter.

## Applying an Enterprise Installation Standards Patch Baseline

The enterprise installation standards (EIS) set of patches is based on the recommended patch cluster for the Oracle Solaris OS. The EIS patch baseline has additional patches included by Oracle field engineers for additional products and to address issues that do not meet the criteria for inclusion in the recommended patch cluster.

## Downloading a Patch

The patches can be downloaded from the [My Oracle Support](#) (MOS) web site.

### ▼ How to Search for a Patch

Before downloading the patch, you might have to search for the patch. To search for a patch, perform the following steps:

- 1 Log in to the MOS web site.
- 2 Click the Patches & Updates tab.
- 3 In the Patch Search section, click Product or the Family (Advanced Search) option.
- 4 Select Solaris Operating System as the Product.
- 5 Select Solaris 10 Operating System as the Release.
- 6 Select Patch or Patchset or both as the type.  
The term Patch refers to a single patch and the term Patchset refers to collection of patches. Patch bundles and patch clusters fall under the category of Patchset.
- 7 Click Search.



---

## Displaying Information About Patches

Before applying patches, you might want to know more information about patches that have been previously applied.

The following commands provide useful information about patches that are already applied to a system:

- `patchadd -p`  
Shows all patches that have been applied to the system.
- `pkgparam pkgid PATCHLIST`  
Shows all patches that have been applied to the package identified by *pkgid*, for example, `SUNWadmap`.

The following `patchadd` command examples to display information about patches that have been applied to your system.

- To obtain information about all patches that have been applied to your system, type:  

```
$ patchadd -p
```
- To verify whether a particular patch has been applied to your system, use the `grep`. For example, to check for a patch with the ID 111879, type:  

```
$ patchadd -p | grep 111879
```

## Applying a Patch

To apply a patch, use the `patchadd` command. For more information about the `patchadd` command, see the [patchadd\(1M\)](#) man page.

---

**Note** – Improvements have been made to the `patchadd -M` command. When you use this command to apply patches to your system, you are no longer required to specify patch IDs in the correct installation order. If you use the `patchadd -M` command without specifying a patch ID, all patches in the directory are installed on the system.

---

The `patchadd` command cannot apply a patch or software update under the following conditions:

- The package is not fully installed on the system.
- The patch package's architecture differs from the system's architecture.
- The patch package's version does not match the installed package's version.
- A patch with the same base code and a higher revision number has already been applied.
- A patch that obsoletes this patch has already been applied.

- The patch is incompatible with a patch that has already been applied to the system.
- The patch being applied depends on another patch that has not yet been applied.

## ▼ How to Apply a Patch Using the `patchadd` Command

Assume that you have downloaded a patch file (`119784-17.zip`) from the MOS web site. To apply the patch to the Oracle Solaris OS, perform the following procedure:

**1 Become superuser.**

**2 Copy the patch file to a temporary directory.**

```
cp /<patch download location>/119784-17.zip /tmp
```

**3 Unzip the patch file.**

```
cd /tmp
unzip 119784-17.zip
```

**4 Apply the patch.**

```
patchadd 119784-17
```

**5 (Optional) Verify that the patch was applied.**

```
patchadd -p | grep 119784-17
```

## Removing a Patch

To remove a patch, use the `patchrm` command. For example to remove the 119784-17 patch, perform the following steps:

1. Become superuser.
2. Remove the patch.

```
patchrm 119784-17
```

To verify that the patch was removed, you can run the `patchadd` command with the `-p` option and search for the specific patch Id. The command should not return any results. For example:

```
patchadd -p | grep 119784-17
```

---

# Patch Management Terms and Definitions

The following terms are used throughout the patch management chapters.

|                              |                                                                                                                                                                                                                                                                                                                           |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>apply</b>                 | To install a patch on a system.                                                                                                                                                                                                                                                                                           |
| <b>back out</b>              | To remove a patch from a system.                                                                                                                                                                                                                                                                                          |
| <b>backout data</b>          | Data that is created when a patch is applied to enable the system to return to its previous state if the patch is removed (backed out).                                                                                                                                                                                   |
| <b>backout directory</b>     | Directory in which backout data is stored. By default, this is the save directory of each package that was installed by the patch.                                                                                                                                                                                        |
| <b>dependency</b>            | See <b>patch dependency</b> .                                                                                                                                                                                                                                                                                             |
| <b>digital signature</b>     | An electronic signature that can be used to ensure that a document has not been modified since the signature was applied.                                                                                                                                                                                                 |
| <b>download</b>              | To copy one or more patches from a source of patches to the system where the patches are to be applied.                                                                                                                                                                                                                   |
| <b>download directory</b>    | Directory in which patches are stored when they are downloaded from the patch source. This is also the directory from which patches are applied. The default location is <code>/var/sadm/spool</code> .                                                                                                                   |
| <b>keystore</b>              | A repository of certificates and keys that is queried when you attempt to apply a signed patch.                                                                                                                                                                                                                           |
| <b>nonstandard patch</b>     | A patch that cannot be installed by using the <code>patchadd</code> command. A nonstandard patch is not delivered in package format. This patch must be installed according to the instructions specified in the README file of the patch. A nonstandard patch typically delivers firmware or application software fixes. |
| <b>order</b>                 | To sort a set of patches in an order suitable for applying patches.                                                                                                                                                                                                                                                       |
| <b>package</b>               | The form in which software products are delivered for installation on a system. The package contains a collection of files and directories in a defined format.                                                                                                                                                           |
| <b>patch</b>                 | An update to software that corrects an existing problem or that introduces a feature.                                                                                                                                                                                                                                     |
| <b>patch analysis</b>        | A method of checking a system to determine which patches are appropriate for the system.                                                                                                                                                                                                                                  |
| <b>patch dependency</b>      | An instance where a patch depends on the existence of another patch on a system. A patch that depends on one or more patches can only be applied to a system when those other patches have already been applied.                                                                                                          |
| <b>patch ID</b>              | A unique alphanumeric string, with the patch base code first, a hyphen, and a number that represents the patch revision number.                                                                                                                                                                                           |
| <b>patch incompatibility</b> | A rare situation where two patches cannot be on the same system. Each patch in the relationship is incompatible with the other. If you want to apply a patch that is incompatible with a patch already on the system, you must first remove the patch that is already on the system. Then, you can apply the new patch.   |
| <b>patch list</b>            | A file that contains a list of patches, one patch ID per line. Such a list can be used to perform patch operations. The list can be generated based on the analysis of a system or on user input.                                                                                                                         |

Each line in a patch list has two columns. The first column is the patch ID, and the second column is a synopsis of that patch.

- patch server** A source of patches that can be used by your systems to perform patch analyses and from which they obtain the appropriate patches.
- software update** A change to software that you apply that corrects an existing problem or that introduces a feature.
- special handling** Patches with properties that indicate they must be installed in single-user mode. Also, patches that require you to restart the system after the patch has been applied are referred to as having *special handling requirements*.
- standard patch** Standard patches are those that adhere to the Oracle Solaris patch specification and are installable by using the `patchadd` command. Note that nonstandard patches cannot be installed by using the `patchadd` command.
- superseded patch** An instance where a patch replaces another patch even if it has not already been applied to a system. For example, if a patch X is marked as superseded, then there exists a patch Y that contains the fixes of patch X. In this example, patch Y replaces the superseded patch X completely. You do not need to apply patch X unless patch X is required for a specific purpose.

## SMF Services

---

The following table lists some of the services that have been converted to use SMF. Each service includes the daemon or service name, the FMRI for that service, the run script that is used to start the service, and whether the service is started by `inetd`.

TABLE A-1 SMF Services

| Service Name | FMRI                                               | Run Script | inetd Service  |
|--------------|----------------------------------------------------|------------|----------------|
| automount    | svc:/system/filesystem/autofs:default              | autofs     | Not applicable |
| consadmd     | svc:/system/consadm:default                        | rootusr    | Not applicable |
| coreadm      | svc:/system/coreadm:default                        | coreadm    | Not applicable |
| cron         | svc:/system/cron:default                           | cron       | Not applicable |
| cryptoadm    | svc:/system/cryptosvc:default                      | N/A        | Not applicable |
| cvcd         | svc:/system/cvc:default                            | cvcd       | Not applicable |
| dcs          | svc:/platform/<arch>/dcs:default                   | None       | Applicable     |
| dtlogin      | svc:/application/graphical-login/cde-login:default | dtlogin    | Not applicable |
| dtprintinfo  | svc:/application/cde-printinfo:default             | dtlogin    | Not applicable |
| dtspcd       | svc:/network/cde-spc:default                       | None       | Applicable     |
| dumpadm      | svc:/system/dumpadm:default                        | savecore   | Not applicable |
| efdaemon     | svc:/platform/<arch>/efdaemon:default              | efcode     | Not applicable |
| fmd          | svc:/system/fmd:default                            | N/A        | Not applicable |
| gssd         | svc:/network/rpc/gss:default                       | None       | Applicable     |

TABLE A-1 SMF Services (Continued)

| Service Name | FMRI                             | Run Script  | inetd Service  |
|--------------|----------------------------------|-------------|----------------|
| imapd        | svc:/network/imap/tcp:default    | None        | Applicable     |
|              | svc:/network/imapnew/tcp:default |             |                |
| in.chargend  | svc:/network/chargen:dgram       | None        | Applicable     |
|              | svc:/network/chargen:stream      |             |                |
| in.comsat    | svc:/network/comsat:default      | None        | Applicable     |
| in.daytimed  | svc:/network/daytime:dgram       | None        | Applicable     |
|              | svc:/network/daytime:stream      |             |                |
| in.dhcpd     | svc:/network/dhcp-server:default | dhcp        | Not applicable |
| in.discardd  | svc:/network/discard:dgram       | None        | Applicable     |
|              | svc:/network/discard:stream      |             |                |
| in.echod     | svc:/network/echo:dgram          | None        | Applicable     |
|              | svc:/network/echo:stream         |             |                |
| in.fingerd   | svc:/network/finger:default      | None        | Applicable     |
| in.ftpd      | svc:/network/ftp:default         | None        | Applicable     |
| in.named     | svc:/network/dns/server:default  | inetsvc     | Not applicable |
| in.rarpd     | svc:/network/rarp:default        | boot.server | Not applicable |
| in.rdisc     | svc:/network/initial:default     | inetinit    | Not applicable |
| in.rexecd    | svc:/network/rexec:default       | None        | Applicable     |
| in.rlogind   | svc:/network/login:rlogin        | None        | Applicable     |
|              | svc:/network/login:eklogin       |             |                |
|              | svc:/network/login:klogin        |             |                |
| in.routed    | svc:/network/initial:default     | inetinit    | Not applicable |
| in.rshd      | svc:/network/shell:default       | None        | Applicable     |
|              | svc:/network/kshell              |             |                |
| in.talkd     | svc:/network/talk:default        | None        | Applicable     |
| in.telnetd   | svc:/network/telnet:default      | None        | Applicable     |
| in.tftpd     | svc:/network/tftp/udp6:default   | None        | Applicable     |

TABLE A-1 SMF Services (Continued)

| Service Name       | FMRI                                                                                                                                                        | Run Script                             | inetd Service  |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------|----------------|
| in.timed           | svc:/network/time:dgram<br>svc:/network/time:stream                                                                                                         | None                                   | Applicable     |
| in.tnamed          | svc:/network/tname:default                                                                                                                                  | None                                   | Applicable     |
| in.uucpd           | svc:/network/uucp:default                                                                                                                                   | None                                   | Applicable     |
| inetd-upgrade      | svc:/network/inetd-upgrade:default                                                                                                                          | N/A                                    | Not applicable |
| inetd              | svc:/network/inetd:default                                                                                                                                  | inetsvc                                | Not applicable |
| intrd              | svc:/system/intrd:default                                                                                                                                   | None                                   | Not applicable |
| ipop3d             | svc:/network/pop3/tcp:default                                                                                                                               | None                                   | Applicable     |
| kadmind            | svc:/network/security/kadmin:default                                                                                                                        | kdc.master                             | Not applicable |
| kbd                | svc:/system/keymap:default                                                                                                                                  | keymap                                 | Not applicable |
| keyserv            | svc:/network/rpc/keyserv:default                                                                                                                            | rpc                                    | Not applicable |
| kpropd             | svc:/network/security/krb5_prop:default                                                                                                                     | None                                   | Applicable     |
| krb5kdc            | svc:/network/security/krb5kdc:default                                                                                                                       | kdc                                    | Not applicable |
| ktkt_warnd         | svc:/network/security/ktkt_warn:default                                                                                                                     | None                                   | Applicable     |
| ldap_cachemgr      | svc:/network/ldap/client:default                                                                                                                            | ldap.client                            | Not applicable |
| loadkeys           | svc:/system/keymap:default                                                                                                                                  | keymap                                 | Not applicable |
| lockd              | svc:/network/nfs/client:default<br>svc:/network/nfs/server:default                                                                                          | nfs.server                             | Not applicable |
| lpsched and lpshut | svc:/application/print/server:default                                                                                                                       | lp                                     | Not applicable |
| mdmonitord         | svc:/system/mdmonitor:default                                                                                                                               | svm.sync                               | Not applicable |
| metainit           | svc:/system/metainit:default                                                                                                                                | svm.init                               | Not applicable |
| metadevadm         | svc:/platform/<arch>/mpxio-upgrade:default                                                                                                                  | N/A                                    | Not applicable |
| mount              | svc:/system/filesystem/local:default<br>svc:/system/filesystem/minimal:default<br>svc:/system/filesystem/root:default<br>svc:/system/filesystem/usr:default | nfs.client, rootusr,<br>standardmounts | Not applicable |
| mountd             | svc:/network/nfs/server:default                                                                                                                             | nfs.server                             | Not applicable |
| nfsd               | svc:/network/nfs/server:default                                                                                                                             | nfs.server                             | Not applicable |

TABLE A-1 SMF Services (Continued)

| Service Name                        | FMRI                                   | Run Script  | inetd Service  |
|-------------------------------------|----------------------------------------|-------------|----------------|
| nfsmapid                            | svc:/network/nfs/client:default        | nfs.server  | Not applicable |
|                                     | svc:/network/nfs/server:default        |             |                |
| nis_cachemgr                        | svc:/network/rpc/nisplus:default       | rpc         | Not applicable |
| nscd                                | svc:/system/name-service-cache:default | nscd        | Not applicable |
| ntpdate                             | svc:/network/ntp:default               | xntpd       | Not applicable |
| ocfserv                             | svc:/network/rpc/ocfserv:default       | ocfserv     | Not applicable |
| picld                               | svc:/system/picl:default               | picld       | Not applicable |
| pmconfig                            | svc:/system/power:default              | power       | Not applicable |
| printd                              | svc:/application/print/cleanup:default | spc         | Not applicable |
| quotaon                             | svc:/system/filesystem/local:default   | ufs_quota   | Not applicable |
| rcapd                               | svc:/system/rcap:default               | rcapd       | Not applicable |
| rpc.bootparamd                      | svc:/network/rpc/bootparams:default    | boot.server | Not applicable |
| rpc.mdcomm                          | svc:/network/rpc/mdcomm:default        | None        | Applicable     |
| rpc.metad                           | svc:/network/rpc/meta:default          | None        | Applicable     |
| rpc.metamedd                        | svc:/network/rpc/metamed:default       | None        | Applicable     |
| rpc.metamhd                         | svc:/network/rpc/metamh:default        | None        | Applicable     |
| rpc.nisd                            | svc:/network/rpc/nisplus:default       | rpc         | Not applicable |
| rpc.nispasswdd                      | svc:/network/rpc/nisplus:default       | rpc         | Not applicable |
| rpc.rexd                            | svc:/network/rpc/rex:default           | None        | Applicable     |
| rpc.rstatd                          | svc:/network/rpc/rstat:default         | None        | Applicable     |
| rpc.rusersd                         | svc:/network/rpc/rusers:default        | None        | Applicable     |
| rpc.smsserverd                      | svc:/network/rpc/smsserver:default     | None        | Applicable     |
| rpc.sprayd                          | svc:/network/rpc/spray:default         | None        | Applicable     |
| rpc.ttdbserverd                     | svc:/network/rpc/ttdbserver:tcp        | None        | Applicable     |
| rpc.walld                           | svc:/network/rpc/wall:default          | None        | Applicable     |
| rpc.yppasswdd and<br>rpc.yppupdated | svc:/network/nis/server:default        | rpc         | Not applicable |
| rquotad                             | svc:/network/nfs/rquota:default        | None        | Applicable     |



TABLE A-1 SMF Services (Continued)

| Service Name                             | FMRI                                                               | Run Script | inetd Service  |
|------------------------------------------|--------------------------------------------------------------------|------------|----------------|
| sadc                                     | svc:/system/sar:default                                            | perf       | Not applicable |
| savecore                                 | svc:/system/dumpadm:default                                        | savecore   | Not applicable |
| sendmail                                 | svc:/network/smtp:sendmail                                         | sendmail   | Not applicable |
| sf880drd                                 | svc:/platform/<arch>/sf880drd:default                              | sf880dr    | Not applicable |
| slpd                                     | svc:/network/slp:default                                           | slpd       | Not applicable |
| sshd                                     | svc:/network/ssh:default                                           | sshd       | Not applicable |
| statd                                    | svc:/network/nfs/client:default<br>svc:/network/nfs/server:default | nfs.server | Not applicable |
| svc.startd                               | svc:/system/svc/restarter:default                                  | N/A        | Not applicable |
| syseventd                                | svc:/system/sysevent:default                                       | devfsadm   | Not applicable |
| sysidpm, sysidns,<br>sysidroot, sysidsys | svc:/system/sysidtool:system                                       | sysid.sys  | Not applicable |
| sysidnet                                 | svc:/system/sysidtool:net                                          | sysid.net  | Not applicable |
| syslogd                                  | svc:/system/system-log:default                                     | syslog     | Not applicable |
| ttymon                                   | svc:/system/console-login:default                                  | inittab    | Not applicable |
| utmpd                                    | svc:/system/utmp:default                                           | utmpd      | Not applicable |
| vold                                     | svc:/system/filesystem/volfs:default                               | volmgt     | Not applicable |
| xntpd                                    | svc:/network/ntp:default                                           | xntpd      | Not applicable |
| ypbind                                   | svc:/network/nis/client:default                                    | rpc        | Not applicable |
| ypserv                                   | svc:/network/nis/server:default                                    | rpc        | Not applicable |
| ypxfrd                                   | svc:/network/nis/server:default                                    | rpc        | Not applicable |
| zoneadm                                  | svc:/system/zones:default                                          | N/A        | Not applicable |
| None                                     | svc:/network/loopback:default                                      | network    | Not applicable |
| None                                     | svc:/network/physical:default                                      | network    | Not applicable |



# Index

---

## Numbers and Symbols

\$ZFS-BOOTFS, ZFS boot options, 192

## A

account information for My Oracle Support, how to obtain, 326

account information for Sun Online account, how to obtain, 326

active menu.lst file, location of, 219

adding

- a package, example of, 420

- a package from a mounted CD (example of), 419

- diskless client OS services (how to), 150

- multiple versions of a package, 392

- packages (prerequisites), 391

- packages from a spool directory (example of), 423

- packages from remote package server (example of), 420

- packages to a spool directory (example of), 425

- packages with administration files, 393

- run control script (how to), 377

- user initialization files, 102

adding and removing packages and patches  
restrictions on

- zones, 392

adding GRUB menu entries, findroot

- command, 224–225

adding missing ARCH=all packages (example of), 165–166

administering GRUB, reference, 188–189

administration file, keyword, 392

aging user passwords, 94, 126, 127

aliases, user login names vs., 87

appliances, definition, 137

application access, Oracle Java Web Console, 79

application access to remote systems, Oracle Java Web Console, 80

application privileges, Oracle Java Web Console, 79

ARCH=all packages, how to add missing, diskless client troubleshooting, 163–170

archive

- booting the failsafe archive

  - GRUB, 256–257

audit events, Oracle Java Web Console, 66

auditing implementation, Oracle Java Web Console, 65

authenticating by specifying named support credentials, Oracle Solaris Auto Registration, 338

authenticating with My Oracle Support  
specifying named support credentials

- auto registration, 326

authentication failed, troubleshooting issues with auto registration, 345

authTypes tag, Oracle Java Web Console, 80

auto registration

- clearing configuration information, 335

- configuration information, 334

- configuring properties, 336

Auto Registration

- data transmittal

  - service tags, 327

- auto registration
    - dry run function, 341
    - enabling SMF services, 332
    - how it works, 328
    - how to obtain information, 344
    - HTTP proxy and port
      - configuring, 336
    - HTTP proxy authentication, 337
    - providing named support credentials, 338
    - providing support credentials, 326
    - regadm command, 330
      - requirements, 331
    - registering anonymously
      - example, 344
    - troubleshooting issues with, 345
    - using the Oracle Solaris Auto Registration feature, 326
  - auto registration CLI, regadm command, 342
  - auto registration issues
    - authentication failed
      - error messages, 345
  - automounting, user home directories, 92
- B**
- banner command (PROM), 205
  - base directory (basedir), 392, 394
  - basedir keyword (administration files), 392
  - bin group, 88
  - BIOS
    - system BIOS
      - in GRUB boot environment, 289–290
  - boot archive, how to rebuild a corrupt, 258–260
  - boot archives
    - managing, 267–276
    - types of, 187–188
  - boot archives, managing, 267–268
  - boot behavior
    - editing the GRUB menu.lst file
      - how to, 219–222
    - how to modify in GRUB menu, 217–218, 253–255
    - managing, 203–225
  - boot behavior, modifying on x86 based systems, 212–225
  - boot-file property, changing, 209–210
  - boot options
    - L
      - ZFS root file system, 234–238
    - Z
      - ZFS root file system, 234–238
  - boot process, x86, 321
  - boot-time interactions, GRUB menu, 218–219
  - bootadm command, using to manage boot archives, 269–270
  - bootfs pool property, 191
  - booting
    - 64-bit x86 based system in 32-bit mode (example of), 312
    - a diskless client (how to), 158
    - a system, guidelines, 178–179
    - interactively (how to)
      - SPARC, 231
    - the Solaris Device Configuration Assistant (how to)
      - x86, 308
    - to run level S
      - SPARC, 230
  - booting a system by using GRUB, overview, 295–296
  - booting a system to run level S
    - GRUB based booting
      - how to, 247–249
  - booting an x86 based system interactively with GRUB, 249–251
  - booting from a ZFS root file system
    - SPARC boot options, 191–192
    - x86 boot options, 192
  - booting from the network with GRUB, 260–264
    - DHCP configuration, 261
  - booting from ZFS root file system, SPARC platform, 234–238
  - booting in failsafe mode
    - GRUB based booting, 256–257
    - SPARC based systems, 239–242
  - booting the failsafe archive, to rebuild a corrupt boot archive, 258–260
  - booting with GRUB, reference, 188–189
  - Bourne shell
    - See also* user initialization files
    - basic features, 104

Break key, 278

## C

### C shell

- basic features, 104
- environment variables and, 105, 109
- shell (local) variables and, 105
- user initialization files and, 102, 110
  - See* user initialization files
- creating, 104

cannot send HTTP proxy password  
 auto registration troubleshooting  
 error messages, 345

### CD-ROM devices

- adding software from mounted CD
- example of, 419

CDPATH environment variable, 106

### changing

- directory ownership for user accounts, 124
- file ownership for user accounts, 124
- Oracle Java Web Console properties
  - session timeout period, 67
- user ID numbers, 123
- user login names, 123
- user passwords
  - by user, 91
  - frequency of, 91
  - Users Tool, 126

changing boot properties, 209–210

changing Oracle Java Web Console properties,  
 choosing an auditing implementation, 65

checking, installed packages (example of), 425

class macro, configuring DHCP, 262

clean shutdown, 194

clearing auto registration configuration, 335

client macro, configuring DHCP, 262

commands (SMF), list of, 354

commands used to register your system

- regadm register
  - auto registration, 342

compatibility with other applications, Oracle Java Web  
 Console, 57

components of GRUB, 292–293

configuration repository (SMF), *See* repository  
 configuring auto registration properties, 336  
 configuring DHCP, booting from the network with  
 GRUB, 261–262

configuring Oracle Java Web Console, 64

console access, Oracle Java Web Console, 79

console session timeout, changing Oracle Java Web  
 Console properties, 64

controlling file and directory access, 109

corrupt boot archive, how to rebuild, 258–260

.cshrc file

- customizing, 104, 110
- description, 102

## D

### -d option

- regadm register
  - obtain registration information in advance, 341

daemon group, 88

data transmittal, service tags, 327

definitions of patch-related terms, 435–436

delegated restarters (SMF), 355

### deleting

- diskless client OS services (example of), 159
- diskless client OS services (how to), 159
- user home directories, 124
- user mailboxes, 124

dependency statements (SMF), description, 348

determining, system's run level (how to), 357

device naming conventions, in GRUB, 292–293

dfstab file, user home directory sharing and, 121

DHCP, configuring a GRUB based PXE boot, 261

DHCP macros, using in GRUB, 261–262

### directories

- base directory (basedir), 392
- changing ownership for user accounts, 124
- controlling access to, 109
- home, 91
- PATH environment variable and, 106, 107
- skeleton, 102

### disabling

- run control script (how to), 378

disabling (*Continued*)

- user accounts
  - passwords and, 124
  - Users tool, 124
- diskless client management commands
  - smossservice
    - add OS services, 141
- diskless client troubleshooting, how to add missing
  - ARCH=all packages, 163–170
- diskless clients
  - adding OS services for (how to), 150
  - booting (how to), 158
  - definition, 136
  - deleting OS services (example of), 159
  - deleting OS services (how to), 159
- displaying
  - environment variables, 105
  - installed software information, 423
  - user mask, 109
- displaying a list of available BEs
  - booting a ZFS root
    - boot -L option, 191–192
- displaying auto registration configuration, 334
- dry run option, regadm register -d, 341

**E**

- editing the menu.lst file, modifying boot
  - behavior, 219–222
- eeprom command
  - how to use to set boot parameters
    - GRUB, 214
    - modifying boot behavior, 213–214
- enabling auto registration, SMF services, 332
- encryption, 94
- env command, 105
- environment variables
  - description, 105, 109
  - LOGNAME, 106
  - LPDEST, 106
  - PATH, 106
  - SHELL, 107
  - TZ, 107

## error messages

- authentication failed
  - auto registration, 345
  - troubleshooting issues with auto registration
    - HTTP communications error, 345
    - HTTP proxy password, 345
- /etc/dfs/dfsstab file, user home directory sharing
  - and, 121
- /etc files
  - user account information and, 92
- /etc/init.d directory, 377
- /etc/inittab file
  - entry description, 358
  - example of default, 359
- /etc/passwd file
  - description, 94
  - fields in, 94
  - user ID number assignment and, 88
  - recovering
    - SPARC, 281
    - recovering (example of)
      - x86, 259, 310
    - deleting user accounts and, 124
- /etc/shadow file, description, 94
- /etc/vfstab file, 122
- /export/home file system, 91

**F**

- failsafe
  - GRUB based booting
    - recovery, 256–257
- failsafe archive, booting, 267–268
- failsafe mode, booting on SPARC based systems, 239–242
- fault management resource identifier, *See* FMRI
- files
  - changing ownership for user accounts, 124
  - controlling access to, 109
  - verifying attributes for newly installed packages, 425
- findroot command
  - adding GRUB menu entries, 224–225
  - menu.lst entries, 223–224

FMRI, description, 350–351  
 forget root password, SPARC, 282  
 functional components of GRUB, 292–293

## G

GIDs, 88  
   assigning, 90  
   definition, 89  
   large, 88  
 glossary of patch-related terms, 435–436  
 group file  
   deleting user accounts and, 124  
   description, 94  
   fields in, 96  
 group ID numbers, 88, 89, 90  
 groups  
   changing primary, 90  
   default, 90  
   description, 89  
   description of names, 89  
   displaying groups a user belongs to, 90  
   guidelines for managing, 89, 90  
   ID numbers, 88, 89, 90  
   name services and, 90  
   names  
     description, 89  
     primary, 90  
     secondary, 90  
   storage of information for, 94, 96  
   UNIX, 89  
 groups command, 90  
 GRUB  
   modifying boot behavior  
     editing the menu.lst file, 219–222  
   support for multiple operating systems, 293–294  
 GRUB based booting  
   about DHCP macros, 261–262  
   booting a system interactively, 249–251  
   booting in failsafe mode, 256–257  
   how to boot a system run level S, 247–249  
   how to rebuild a corrupt boot archive, 258–260  
   modifying kernel Behavior in the GRUB menu at  
     boot time, 214–215

GRUB based booting (*Continued*)  
   modifying the GRUB kernel behavior at boot  
     time, 217–218, 253–255  
 GRUB-based booting, reference, 188–189  
 GRUB based network boot, 260–264  
 GRUB device naming conventions, 292–293  
 GRUB functional components, 292–293  
 GRUB menu  
   description of, 218–219  
   modifying GRUB kernel behavior, 214–215  
 GRUB terminology, 290–291  
 GRUBClient, GRUB based network boot, 260–264

## H

halt command, 195  
 history environment variable, 106  
 HOME environment variable, 106  
 /home file system, user home directories and, 91  
 how auto registration works, Oracle Solaris Auto  
   Registration, 328  
 how products are registered, Oracle Solaris Auto  
   Registration, 326  
 how to use GRUB to boot a system to run level  
   s, 247–249  
 HTTP communications error, auto registration  
   troubleshooting, 345  
 HTTP proxy and port settings  
   auto registration configuration  
     configuring, 336  
 HTTP proxy authentication, auto registration  
   configuration, 337  
 HTTP proxy password  
   error messages  
     troubleshooting auto registration, 345

## I

ID numbers  
   group, 88, 89, 90  
   user, 88, 123  
 inetadm command, description, 354

- init command
  - description, 195
  - shutting down a stand-alone system, 200
- init states, *See* run levels
- initialization files, system, 92
- interactive boot, booting an x86 based system with GRUB, 249–251
- inventory
  - registering by using the Auto Registration feature, 326–330
  - service tags, 344
- IP macro, configuring DHCP, 262

## K

- kernel behavior, modifying in GRUB menu, 214–215
- kernel initialization in the GRUB boot environment, 290
- Korn shell
  - basic features, 104
  - user initialization files and, 102

## L

- L1-A keys, 278
- L boot option, booting a ZFS root file system on SPARC platform, 234–238
- L option
  - ZFS boot options
    - displaying available BEs, 191–192
- LANG environment variable, 106, 108, 109
- LC environment variables, 108, 109
- legacy applications, Oracle Java Web Console, 74
- library interfaces, SMF, 354
- listing, package information (example of), 423
- \*LK\* password, 124
- local.cshrc file, 102
- local.login file, 102
- local.profile file, 102
- locale environment variable, 106
- location of active menu.lst file, 219
- .login file
  - customizing, 104, 110

- .login file (*Continued*)
  - description, 102
- login names (user)
  - changing, 123
  - description, 87
- LOGNAME environment variable, 106
- LPDEST environment variable, 106

## M

- macros, DHCP, 261–262
- mail aliases, user login names vs., 87
- MAIL environment variable, 105, 106
- managing boot behavior, 203–225
- managing Oracle Java Web Console service, 62–64
- managing services, for auto registration, 331
- managing the boot - archive service, 269–270
- managing the boot archives, tasks, 267–276
- manifests (SMF), description, 351
- MANPATH environment variable, 106
- maximums
  - secondary groups users can belong to, 90
  - user ID number, 88
  - user login name length, 93
  - user password length, 91
- menu
  - GRUB
    - description of, 218–219
- menu.lst, GRUB component, 292–293
- menu.lst file
  - adding entries that use the findroot command, 224–225
  - and boot-time interactions
    - description, 218–219
  - location, 219
  - modifying boot behavior, 219–222
  - multiboot implementation, 296
- minimums
  - user login name length, 93
  - user password length, 91
- modifying boot behavior
  - editing the GRUB menu.lst file
    - how to, 219–222
  - in GRUB menu at boot time, 213–214



modifying boot behavior (Task Map), 212–225  
 modifying kernel usage in the GRUB menu, 217–218,  
 253–255  
 mounting  
   user home directories  
     automounting, 92  
   user home directories (how to), 122  
 multiboot implementation, menu.lst file  
   description, 296  
 multiple operating systems in the GRUB boot  
   environment, 293–294  
 multiple versions of software packages, 392, 394  
 multiuser level, *See* run level 3  
 My Oracle Support, account information, 326

## N

name services  
   groups and, 90  
   user accounts and, 92, 94  
 names  
   group  
     description, 89  
   software package naming conventions, 391  
   SUNW prefix, 391  
   user login  
     changing, 123  
     description, 87  
 naming conventions for devices, in GRUB, 292–293  
 Navigation pane of Solaris Management Console,  
   nodes, 35  
 network boot, with GRUB, 260–264  
 network macro, configuring DHCP, 262  
 new features, SMF, 347  
 newgrp command, 90  
 NIS  
   user accounts and, 92, 94  
 NIS+, user accounts and, 124  
 noaccess user/group, 88  
   and Oracle Java Web Console, 68  
 noask\_pkgadd administration file, 393, 420  
 nobody user/group, 88  
 nodes, Navigation pane of Solaris Management  
   Console, 35

normal archive in GRUB  
   boot archive  
     reference, 187–188  
 notifying users of system down time, 195

## O

obtaining product registration information, how  
   to, 344  
 obtaining registration information in advance  
   dry run function  
     auto registration, 341  
 Oracle Java Web Console, 55  
   (Overview), 56  
   access to applications, 79  
   access to console, 79  
   application access to remote systems, 80  
   application privileges, 79  
   authorizing users of applications, 80  
   changing properties of  
     auditing implementation, 65  
     console session timeout, 64  
     logging level, 64  
   changing the user identity that runs the console, 68  
   compatibility with other applications, 57  
   configuring, 64  
   configuring properties, 66–68  
   differences between default logging and debug  
     logging, 65  
   disabling the console service, 63–64  
   enabling the console service, 62–63  
   internal passwords, 80  
   legacy applications, 74  
   listing deployed applications, 74–75  
   managing the console service, 62–64  
   noaccess user identity, 68  
   properties, 71–73  
   reference information, 78–83  
   registering applications, 75–76, 77  
   security considerations, 78  
   starting applications from, 59  
   starting the console service, 62  
   status, 71–73  
   stopping the console service, 63

- Oracle Java Web Console (*Continued*)
    - troubleshooting, 71
    - unregistering applications, 76–77, 78
    - using `authTypes` tag, 80
  - Oracle Java Web Console commands
    - `smcwebserver`, 58
    - `smreg`, 58
    - `wcadmin`, 58
  - Oracle Product Registration System
    - obtaining information, 344
    - transmitting data, 327
  - Oracle Solaris Auto Registration
    - clearing configuration, 335
    - configuring properties, 336
    - displaying configuration, 334
    - how it works, 328
    - HTTP proxy and port settings
      - configuring, 336
    - HTTP proxy authentication
      - configuring, 337
    - managing settings, 328
    - managing SMF services, 331
    - overview, 326–330
    - using the `regadm register` command, 342
  - Oracle Solaris Auto Registration feature,
    - troubleshooting, 345
  - Oracle Solaris boot behavior, how to manage, 203–225
  - OS server, description, 141
- P**
- packages
    - adding
      - See also* `pkgadd` command
    - definition of, 389
    - overview, 389
  - `passwd` file, 94
    - deleting user accounts and, 124
    - fields in, 94
    - recovering
      - SPARC, 281
    - recovering (example of)
      - x86, 259, 310
    - user ID number assignment and, 88
  - passwords (user)
    - aging, 94, 126, 127
    - changing
      - frequency of, 91
      - by user, 91
      - Users Tool, 126
    - choosing, 91
    - description, 91, 127
    - disabling/locking user accounts and, 124
    - encryption, 94
    - \*LK\* password, 124
    - precautions, 91
    - setting, 91, 126
    - Users Tool, 126
  - patches, terms used with, 435–436
  - `PATH` environment variable
    - description, 106, 107
  - `path` shell variable, 105
  - permissions, 109
  - `/pkg` directory, 423
  - `pkgadd` command
    - `-d` option (device name), 418, 421, 422
    - `-s` option (spool directory), 421, 422
    - adding packages (how to), 418
      - using an HTTP URL, 421
    - alternate base directory and, 394
    - bypassing user interaction, 393, 394
    - overview, 390
    - `-a` option (administration file), 393, 394, 418, 420
    - prerequisites for using, 391
    - spool directories and, 421
    - spool directories and (example of), 423
  - `pkgadm` command, overview, 390
  - `pkgchk` command
    - overview, 390
    - using (example of), 425
  - `pkginfo` command
    - displaying all packages installed (example of), 423
    - how to use, 423
    - overview, 390, 391
  - `pkgparam` command, overview, 390
  - `pkgrm` command
    - caution, 392
    - overview, 390

- pkgrm command (*Continued*)
    - prerequisites for using, 391
    - rm command (compared), 392
  - pkgtrans command, overview, 390
  - poweroff command, 195
  - primary administrator role
    - creating (overview), 41
    - granting rights, 41
  - primary groups, 90
  - processes, how auto registration works, 328
  - prodreg command, overview, 390
  - product registration
    - by using the regadm register command, 342
    - how it works
      - auto registration, 328
      - service tag information, 327
  - product registration information, how to obtain, 344
  - .profile file
    - customizing, 104, 110
    - description, 102
  - profiles (SMF), description, 352
  - PROM, finding the PROM revision, 205
  - prompt shell variable, 106
  - properties
    - auto registration configuration, 336
    - changing the boot-file property, 209–210
  - proxy and port settings
    - auto registration configuration
      - configuring, 336
  - proxy authentication
    - auto registration configuration
      - configuring, 337
  - PS1 environment variable, 106
  - pseudo-ttys, 88
  - pseudo user logins, 88
  - PXEClient, GRUB based network boot, 260–264
- R**
- reboot command, 195
  - rebuilding corrupt boot archive (how to), 258–260
  - recover root password (how to), SPARC, 282
  - recovering
    - booting the failsafe archive
      - GRUB based booting, 256–257
  - reference, administering GRUB, 188–189
  - regadm command
    - configuring properties, 336
    - requirements for using, 331
    - when to use, 330
  - regadm register, auto registration tasks, 342
  - regadm register -d, obtaining registration information in advance, 341
  - registered products, obtaining information about, 344
  - registering anonymously (example), 344
  - registering system assets, Oracle Solaris Auto Registration, 326–330
  - registering with support credentials (example), 344
  - registering your system
    - how auto registration works, 328
    - using the Oracle Solaris Auto Registration feature, 326
    - using the regadm register command, 342
  - remote package server
    - adding packages to a spool directory (example of), 422
    - software installation from, 420
    - software installation from (example of), 420
  - removef command, 392
  - removing
    - packages with administration files and, 393
    - software packages
      - guidelines for, 392
  - removing and adding packages and patches
    - restrictions on
      - zones, 392
  - repairing the /etc/passwd file
    - SPARC, 281
    - x86, 259, 310
  - repository (SMF)
    - description, 348, 352
  - requirements for using regadm command, auto registration, 331
  - reset command, 209
  - resetting, a SPARC based system, 209
  - restarters (SMF), 355

- restarters (SMF) (*Continued*)
  - description, 348
- restrictions, on adding and removing packages and patches, 392
- root password, forget, SPARC, 282
- run control scripts
  - adding (how to), 377
  - disabling (how to), 378
  - starting and stopping services, 377
- run level
  - 0 (power-down level), 356
  - 1 (single-user level), 356
  - 2 (multiuser level), 356
  - 3 (multiuser wirooth NFS)
    - booting to, 246
  - 3 (multiuser with NFS), 356
    - booting to, 158, 229, 299
    - what happens when system is brought to, 359
  - 6 (reboot level), 357
  - default run level, 356
  - definition, 356
  - determining (how to), 357
  - s or S (single-user level), 356
    - booting to, 302
  - s or S (single-user state)
    - booting to, 230
- S**
- secondary groups, 90
- security, user ID number reuse and, 88
- security considerations, Oracle Java Web Console, 78
- selecting a logging level, changing Oracle Java Web Console properties, 64
- servers, OS server, 141
- service (SMF), description, 349
- service configuration repository, *See* repository
- service management facility
  - See* SMF
- service states, description, 351
- service tags, product registration, 344
- session timeout period, changing Oracle Java Web Console properties, 67
- set command, 105
- setenv command, 105
- setting Auto Registration properties, by using regadm command, 336
- Setting boot behavior by using eeprom command, GRUB based booting, 214
- settings
  - auto registration configuration
    - HTTP proxy and port, 336
- shadow file
  - description, 94
  - fields in, 95
- SHELL environment variable, 107
- shell variables, 106
- shells
  - basic features, 104
  - environment of, 105
  - environment variables and, 105, 109
  - local variables, 105, 106
  - user initialization files and, 110
- shutdown command
  - description, 195
  - notifying users, 195
  - shutting down a server, 178
  - shutting down a server (how to), 196
- shutting down
  - a system, guidelines, 177–178
  - a system cleanly with shutdown and init commands, 194
- single sign-on, secure https port, Java Web Console, 57
- single-user level, *See* run level s or S
- site initialization files, 103
- skeleton directories (/etc/skel), 102
- smcwebserver command, Oracle Java Web Console, 58
- SMF
  - commands, 354
  - delegated restarters, 355
  - library interfaces, 354
  - overview, 347
- SMF services for auto registration,
  - svc:/application/autoreg, 331
- smreg command
  - Oracle Java Web Console, 58, 76

- snapshots (SMF), description, 353
  - software management
    - naming conventions for packages, 391
    - packages and, 389
    - tools for, 390
  - software packages
    - installing, 423
    - installing from a spool directory (example of), 422
  - Solaris Device Configuration Assistant,
    - overview, 307–308
  - Solaris Management Console
    - description, 31
    - description of tools, 32
    - reasons for using, 34
    - starting (how to), 44
    - using with RBAC, 40
  - SPARC boot options, booting from an Oracle Solaris ZFS root file system, 191–192
  - specifying named support credentials, Oracle Solaris Auto Registration, 338
  - spool directories
    - installing software packages to (example of), 423, 425
    - installing software packages to (how to), 421
  - staff group, 90
  - stage2, GRUB component, 292–293
  - stand-alone systems, definition, 136
  - starting and stopping services, 377
  - starting applications, Oracle Java Web Console launch page, 59
  - Stop-A keys, 278
  - stopping
    - a system for recovery purposes
      - SPARC, 278
    - a system for recovery purposes (how to)
      - x86, 284, 308
  - stty command, 108
  - Sun Online account, how to obtain, 326
  - Sun software packages
    - adding (example of), 419
    - installing, 420
  - SUNW prefix, 391
  - superuser (root) password, forget, SPARC, 282
  - support credentials
    - Oracle Solaris Auto Registration
      - providing, 338
    - providing, 326
    - registering by specifying
      - example, 344
      - when to provide, 327
  - svc.startd daemon, description, 354–355
  - svcadm command, description, 354
  - svccfg command, description, 354
  - svccprop command, description, 354
  - svcs command, description, 354
  - sync command, 279
  - synchronize file systems with sync command, 279
  - system accounts, 88
  - system BIOS in GRUB boot environment, 289–290
  - system initialization files, 92
  - system shutdown commands, 195
  - system types
    - appliance, 137
    - diskless client, 136
    - guidelines for choosing, 137
    - overview, 135
    - stand-alone system, 136
- T**
- TERM environment variable, 107
  - TERMINFO environment variable, 107
  - terminology, GRUB, 290–291
  - time zone environment variable, 107
  - transmitting data, service tags, 327
  - troubleshooting
    - a failed 64-bit boot, 315
    - Oracle Java Web Console, 71
  - troubleshooting auto registration issues
    - HTTP communications error, 345
    - HTTP proxy password
      - error messages, 345
  - troubleshooting diskless client installation problems,
    - adding missing ARCH=all packages (how to), 163–170
  - troubleshooting failed authentication, auto registration issues, 345

troubleshooting issues with auto registration, 345  
ttys (pseudo), 88  
ttytype pseudo user logins, 88  
TZ environment variable, 107

## U

UIDs, 123  
    assigning, 88  
    definition, 88  
    large, 88  
umask command, 109  
UNIX groups, 89  
unregistering an application from the Oracle Java Web Console, 76  
user accounts, 86  
    description, 86, 87  
    disabling/locking  
        passwords and, 124  
        Users Tool, 124  
    guidelines for, 92  
    ID numbers, 88, 123  
    login names, 87, 123  
    name services and, 92, 94  
    setting up  
        information sheet, 114  
    storage of information for, 92  
user home directories  
    changing ownership of, 124  
    customized initialization files in, 102  
    deleting, 124  
    description, 91  
    mounting  
        automounting, 92  
        mounting (how to), 122  
        nonlocal reference to (\$HOME), 92, 103  
user ID numbers, 88, 123  
user initialization files  
    customizing, 101, 110  
        adding customized files, 102  
        avoiding local system references, 103  
        environment variables, 105, 109  
        overview, 102  
        shell variables, 106, 107

user initialization files, customizing (*Continued*)  
    site initialization files, 103  
    user mask setting, 109  
description, 92  
examples, 110  
shells and, 103, 104, 110  
user login names  
    changing, 123  
    description, 87  
user logins (pseudo), 88  
user mask, 109  
Users Tool  
    disabling accounts, 124  
    password administration, 126  
uucp group, 88

## V

/var/sadm/install/admin directory, 393  
/var/spool/pkg directory, 421, 423  
variables  
    environment, 105, 109  
    shell (local), 105  
verifying  
    software installation (example of), 425  
    software package installation  
        pkginfo command, 421  
    software package installation with pkginfo  
        command, 421  
versions of GRUB in Oracle Solaris OS, 295–296

## W

wcadmin command, Oracle Java Web Console, 58  
web-based system management applications, Oracle  
    Java Web Console, 56  
who command, 196, 357

## X

x86 boot options, booting from a ZFS root file  
    system, 192

**Z**

- Z boot option, booting a ZFS root file system on SPARC platform, 234–238
- Z option, ZFS boot options, 191–192
- ZFS booting on SPARC platform, boot options used, 234–238
- ZFS root file system, booting from on SPARC platform, 234–238
- zones, restrictions on adding and removing packages and patches, 392

