



# System Administration Guide: Security Services

---

Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054  
U.S.A.

Part No: 806-4078-10  
May 2002

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

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

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

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, SunOS, Java, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. Portions of the software copyright 1996 by the Massachusetts Institute of Technology. All rights reserved.

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

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

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

---

Copyright 2002 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 U.S.A. Tous droits réservés

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

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

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, SunOS, Java, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. Copyright 1996 des portions du logiciel par Massachusetts Institute of Technology. Tous droits réservés.

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

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



020129@3062



# Contents

---

<b>Preface</b>	<b>17</b>
<b>1 Security Services (Overview)</b>	<b>21</b>
Introduction to Security Services	21
Authentication	22
Access Control	22
Secure Communication	23
Auditing	23
<b>2 Authentication Services Topics</b>	<b>25</b>
<b>3 Using Authentication Services (Tasks)</b>	<b>27</b>
Overview of Secure RPC	27
NFS Services and Secure RPC	27
DES Encryption	28
Kerberos Authentication	28
Diffie-Hellman Authentication	28
Administering Diffie-Hellman Authentication	32
▼ How to Restart the Keyserver	32
▼ How to Set Up a Key in NIS+ Credentials for Diffie-Hellman Authentication	32
▼ How to Set Up a New User Key That Uses NIS+ Credentials for Diffie-Hellman Authentication	33
▼ How to Set Up a root Key by Using NIS Credentials With Diffie-Hellman Authentication	34

	▼ How to Create a New User Key That Uses NIS Credentials with Diffie-Hellman Authentication	34
	▼ How to Share and Mount Files With Diffie-Hellman Authentication	35
	PAM (Overview)	36
	Benefits of Using PAM	36
	PAM Components	37
	Stacking Feature	38
	Password-Mapping Feature	38
	Changes to PAM for the Solaris 9 Release	38
	PAM (Tasks)	38
	PAM (Task Map)	39
	Planning for PAM	39
	▼ How to Add a PAM Module	40
	▼ How to Prevent Unauthorized Access From Remote Systems With PAM	41
	▼ How to Initiate PAM Error Reporting	41
	PAM (Reference)	42
	PAM Modules	42
	PAM Module Types	44
	PAM Configuration File	44
<b>4</b>	<b>Using Secure Shell (Tasks)</b>	<b>49</b>
	Introduction to Secure Shell	49
	Using Secure Shell (Task Map)	51
	Using Secure Shell	52
	▼ How to Create a Public/Private Key Pair	52
	▼ How to Log In to Another Host with Secure Shell	53
	▼ How to Log in With No Password While Using <code>ssh-agent</code>	54
	▼ How to Set <code>ssh-agent</code> to Run Automatically	55
	▼ How to Use Secure Shell Port Forwarding	56
	▼ How to Copy Files With Secure Shell	57
	Transferring Files Remotely Using <code>sftp</code>	58
	▼ How to Set Up Default Connections to Hosts Outside a Firewall	59
<b>5</b>	<b>Secure Shell Administration (Reference)</b>	<b>61</b>
	A Typical Secure Shell Session	61
	Session Characteristics	61
	Authentication	62

	Command Execution and Data Forwarding	62
	Configuring the Secure Shell	63
	Secure Shell Client Configuration	63
	Secure Shell Server Configuration	65
	Maintaining Known Hosts on a Site-Wide Basis	67
	Secure Shell Files	67
<b>6</b>	<b>Introduction to SEAM</b>	<b>71</b>
	What Is SEAM?	71
	How SEAM Works	72
	Initial Authentication: the Ticket-Granting Ticket	73
	Subsequent Authentications	75
	Principals	76
	Realms	76
	SEAM Security Services	78
	SEAM Releases	78
	SEAM 1.0 Components	79
	SEAM Components in the Solaris 8 Release	80
	SEAM 1.0.1 Components	80
	SEAM Components in the Solaris 9 Release	81
<b>7</b>	<b>Planning for SEAM</b>	<b>83</b>
	Why Plan for SEAM?	83
	Realms	84
	Realm Names	84
	Number of Realms	84
	Realm Hierarchy	84
	Mapping Host Names Onto Realms	85
	Client and Service Principal Names	85
	Ports for the KDC and Admin Services	86
	Slave KDCs	86
	Database Propagation	87
	Clock Synchronization	87
	Online Help URL	87
<b>8</b>	<b>Configuring SEAM (Tasks)</b>	<b>89</b>
	Configuring SEAM (Task Map)	89

Configuring KDC Servers	90
▼ How to Configure a Master KDC	91
▼ How to Configure a Slave KDC	95
Configuring Cross-Realm Authentication	98
▼ How to Establish Hierarchical Cross-Realm Authentication	98
▼ How to Establish Direct Cross-Realm Authentication	99
Configuring SEAM NFS Servers	101
▼ How to Configure SEAM NFS Servers	101
▼ How to Create a Credential Table	103
▼ How to Add a Single Entry to the Credential Table	103
▼ How to Set Up a Secure NFS Environment With Multiple Kerberos Security Modes	104
Configuring SEAM Clients	106
▼ How to Configure a SEAM Client	106
Setting Up Root Authentication to Mount NFS File Systems	109
Synchronizing Clocks between KDCs and SEAM Clients	109
Swapping a Master KDC and a Slave KDC	111
▼ How to Configure a Swappable Slave KDC	111
▼ How to Swap a Master KDC and a Slave KDC	112
Administering the Kerberos Database	114
Backing Up and Propagating the Kerberos Database	114
▼ How to Back Up the Kerberos Database	116
▼ How to Restore the Kerberos Database	117
▼ How to Manually Propagate the Kerberos Database to the Slave KDCs	118
Setting Up Parallel Propagation	118
▼ How to Set Up Parallel Propagation	119
Administering the Stash File	120
▼ How to Remove a Stash File	120
Increasing Security	120
▼ How to Restrict Access to KDC Servers	120
<b>9 SEAM Error Messages and Troubleshooting</b>	<b>123</b>
SEAM Error Messages	123
SEAM Administration Tool Error Messages	123
Common SEAM Error Messages (A-M)	124
Common SEAM Error Messages (N-Z)	130
SEAM Troubleshooting	133
Problems Mounting a Kerberized NFS File System	133

	Problems Authenticating as root	134
<b>10</b>	<b>Administering Principals and Policies (Tasks)</b>	<b>135</b>
	Ways to Administer Principals and Policies	135
	SEAM Administration Tool	136
	Command-Line Equivalents of the SEAM Tool	137
	Files Modified by the SEAM Tool	137
	Print and Online Help Features of the SEAM Tool	137
	Working With Large Lists in the SEAM Tool	138
	▼ How to Start the SEAM Tool	139
	Administering Principals	140
	Administering Principals (Task Map)	141
	Automating the Creation of New Principals	141
	▼ How to View the List of Principals	142
	▼ How to View a Principal's Attributes	144
	▼ How to Create a New Principal	146
	▼ How to Duplicate a Principal	148
	▼ How to Modify a Principal	148
	▼ How to Delete a Principal	149
	▼ How to Set Up Defaults for Creating New Principals	150
	▼ How to Modify the Kerberos Administration Privileges	151
	Administering Policies	153
	Administering Policies (Task Map)	153
	▼ How to View the List of Policies	154
	▼ How to View a Policy's Attributes	156
	▼ How to Create a New Policy	158
	▼ How to Duplicate a Policy	160
	▼ How to Modify a Policy	160
	▼ How to Delete a Policy	161
	SEAM Tool Reference	162
	SEAM Tool Panel Descriptions	162
	Using the SEAM Tool With Limited Kerberos Administration Privileges	165
	Administering Keytab Files	166
	Administering Keytabs Task Map	167
	▼ How to Add a Service Principal to a Keytab File	168
	▼ How to Remove a Service Principal From a Keytab File	169
	▼ How to Display the Keylist (Principals) in a Keytab File	170

▼ How to Temporarily Disable Authentication for a Service on a Host 171

**11 Using SEAM (Tasks) 173**

Ticket Management 173

Do You Need to Worry About Tickets? 173

▼ How to Create a Ticket 174

▼ How to View Tickets 174

▼ How to Destroy Tickets 176

Password Management 176

Advice on Choosing a Password 177

Changing Your Password 178

**12 SEAM (Reference) 181**

SEAM Files 181

PAM Configuration File 182

SEAM Commands 183

SEAM Daemons 184

SEAM Terminology 184

Kerberos-Specific Terminology 184

Authentication-Specific Terminology 185

Types of Tickets 186

How the Authentication System Works 189

Gaining Access to a Service Using SEAM 190

Obtaining a Credential for the Ticket-Granting Service 190

Obtaining a Credential for a Server 191

Obtaining Access to a Specific Service 192

Using the gsscred Table 193

**13 Managing System Security Topics 195**

**14 Managing System Security (Overview) 197**

Controlling Access to a Computer System 197

Maintaining Physical Site Security 198

Maintaining Login Control 198

Restricting Access to Data in Files 198

Maintaining Network Control 198



	Monitoring System Usage	199
	Setting the Correct Path	199
	Securing Files	199
	Installing a Firewall	200
	Reporting Security Problems	200
	File Security	200
	Commands for File System Security	200
	File Encryption	201
	Access Control Lists (ACLs)	201
	System Security	202
	Login Access Restrictions	202
	Special Logins	203
	Managing Password Information	203
	Using the Restricted Shell	204
	Tracking Superuser (Root) Login	205
	Network Security	206
	Firewall Systems	206
	Authentication and Authorization	208
	Sharing Files	209
	Restricting Superuser (Root) Access	209
	Using Privileged Ports	209
	Using the Automated Security Enhancement Tool (ASET)	210
<b>15</b>	<b>Securing Files (Tasks)</b>	<b>211</b>
	File Security Features	211
	User Classes	211
	File Permissions	212
	Directory Permissions	212
	Special File Permissions (setuid, setgid and Sticky Bit)	213
	Default umask Setting	214
	Displaying File Information	215
	▼ How to Display File Information	215
	Changing File Ownership	217
	▼ How to Change the Owner of a File	217
	▼ How to Change Group Ownership of a File	218
	Changing File Permissions	219
	▼ How to Change Permissions in Absolute Mode	221

▼ How to Change Special Permissions in Absolute Mode	222
▼ How to Change Permissions in Symbolic Mode	223
Searching for Special Permissions	224
▼ How to Find Files With <code>setuid</code> Permissions	224
Executable Stacks and Security	225
▼ How to Disable Programs From Using Executable Stacks	226
▼ How to Disable Executable Stack Message Logging	226
Using Access Control Lists (ACLs)	226
ACL Entries for Files	227
ACL Entries for Directories	228
▼ How to Set an ACL on a File	229
▼ How to Copy an ACL	231
▼ How to Check If a File Has an ACL	231
▼ How to Modify ACL Entries on a File	232
▼ How to Delete ACL Entries From a File	233
▼ How to Display ACL Entries for a File	233
<b>16 Securing Systems (Tasks)</b>	<b>235</b>
Displaying User Login Information	235
▼ How to Display a User's Login Status	235
▼ How to Display Users Without Passwords	236
Temporarily Disabling User Logins	237
Creating the <code>/etc/nologin</code> File	237
▼ How to Temporarily Disable User Logins	237
Saving Failed Login Attempts	238
▼ How to Save Failed Login Attempts	238
Password Protection Using Dial-up Passwords	239
▼ How to Create a Dial-up Password	240
▼ How to Temporarily Disable Dial-up Logins	241
Restricting Superuser ( <code>root</code> ) Access on the Console	241
▼ How to Restrict Superuser ( <code>root</code> ) Login to the Console	242
Monitoring Who Is Using the <code>su</code> Command	242
▼ How to Monitor Who Is Using the <code>su</code> Command	243
▼ How to Display Superuser ( <code>root</code> ) Access Attempts to the Console	243
Modifying a System's Abort Sequence	243
▼ How to Disable or Enable a System's Abort Sequence	244

<b>17</b>	<b>Role-Based Access Control (Overview)</b>	<b>245</b>
	RBAC: Replacing the Superuser Model	245
	Solaris RBAC Elements	246
	Privileged Applications	248
	Applications That Check UIDs and GIDs	249
	Applications That Check Authorizations	249
	Profile Shell	249
	RBAC Roles	250
	RBAC Authorizations	250
	RBAC Rights Profiles	251
	Name Service Scope	251
<b>18</b>	<b>Role-Based Access Control (Tasks)</b>	<b>253</b>
	Configuring RBAC (Task Map)	254
	Planning for RBAC	254
	▼ How to Plan Your RBAC Implementation	254
	First-Time Use of the User Tool Collection	256
	▼ How to Run the User Tool Collection	256
	Setting Up Initial Users	257
	▼ How to Create Initial Users by Using the User Accounts Tool	257
	Setting Up Initial Roles	259
	▼ How to Create the First Role (Primary Administrator) by Using the Administrative Roles Tool	259
	Making <code>root</code> a Role	261
	▼ How to Make <code>root</code> a Role	261
	Managing RBAC Information (Task Map)	262
	Using Privileged Applications	263
	▼ How to Assume a Role at the Command Line	263
	▼ How to Assume a Role in the Console Tools	264
	Creating Roles	265
	▼ How to Create a Role by Using the Administrative Roles Tool	265
	▼ How to Create a Role From the Command Line	267
	Changing Role Properties	268
	▼ How to Change a Role by Using the Administrative Roles Tool	268
	▼ How to Change a Role From the Command Line	270
	Creating or Changing a Rights Profile	270
	▼ How to Create or Change a Rights Profile by Using the Rights Tool	270
	▼ How to Change Rights Profiles From the Command Line	274

Modifying a User's RBAC Properties	274
▼ How to Modify a User's RBAC Properties by Using the User Accounts Tool	275
▼ How to Modify a User's RBAC Properties From the Command Line	275
Securing Legacy Applications	276
▼ How to Add Security Attributes to a Legacy Application	276
▼ How to Add Security Attributes to Commands in a Script	276
▼ How to Check for Authorizations in a Script or Program	276
<b>19 Role-Based Access Control (Reference)</b>	<b>279</b>
RBAC Elements: Reference Information	279
Configuring Recommended Roles	279
Contents of Rights Profiles	280
Authorizations	284
Databases That Support RBAC	285
RBAC Database Relationships	286
The user_attr Database	287
The auth_attr Database	288
The prof_attr Database	290
The exec_attr Database	291
The policy.conf File	292
RBAC Commands	292
Command-Line Applications for Managing RBAC	292
Commands That Require Authorizations	293
<b>20 Using the Automated Security Enhancement Tool (Tasks)</b>	<b>295</b>
Automated Security Enhancement Tool (ASET)	295
ASET Security Levels	296
ASET Tasks	297
ASET Execution Log	299
ASET Reports	300
ASET Master Files	303
ASET Environment File (asetenv)	304
Configuring ASET	304
Restoring System Files Modified by ASET	307
Network Operation With the NFS System	307
ASET Environment Variables	308

	ASET File Examples	311
	Running ASET	313
	▼ How to Run ASET Interactively	313
	▼ How to Run ASET Periodically	314
	▼ How to Stop Running ASET Periodically	315
	▼ How to Collect ASET Reports on a Server	315
	Troubleshooting ASET Problems	316
	ASET Error Messages	316
<b>21</b>	<b>Auditing Topics</b>	<b>321</b>
<b>22</b>	<b>BSM (Overview)</b>	<b>323</b>
	What Is Auditing?	323
	How Does Auditing Work?	324
	How Is Auditing Related to Security?	325
	BSM Terminology	325
	Audit Events	326
	Audit Classes	327
	Audit Flags	327
	Audit Records	328
	Audit Directory	328
	Device Allocation	328
<b>23</b>	<b>Audit Planning</b>	<b>329</b>
	Handling the Audit Trail	329
	Deciding Who and What to Audit	330
	Determining Which Audit Policies to Use	331
	Controlling Auditing Costs	333
	Cost of Increased Processing Time	333
	Cost of Analysis	333
	Cost of Storage	334
	Auditing Efficiently	334
<b>24</b>	<b>Managing the BSM Service (Tasks)</b>	<b>337</b>
	Managing BSM (Task Map)	337
	Configuring Audit Files	338

Configuring Audit Files (Task Map)	338
▼ How to Change Audit Flags	338
▼ How to Change Users' Audit Characteristics	340
▼ How to Change Audit Classes	341
▼ How to Change Audit Events	342
Configuring the Audit Service	343
Configuring the Audit Service (Task Map)	343
▼ How to Create Partitions for Auditing	343
▼ How to Configure the audit_warn Alias	346
▼ How to Enable or Disable an Audit Policy	346
▼ How to Enable Auditing	347
▼ How to Disable Auditing	348
Managing Audit Records	348
Managing Audit Records (Task Map)	349
▼ How to Merge Audit Records	349
▼ How to Display Audit Record Formats	351
▼ How to Prevent Audit Trail Overflow	351
Managing Device Allocation	352
Adding an Allocatable Device (Task Map)	352
▼ How to Set Up Lock Files for an Allocatable Device	353
▼ How to Change Which Devices Can Be Allocated	353
▼ How to Allocate a Device	354
▼ How to Deallocate a Device	354
<b>25 BSM Services Reference</b>	<b>357</b>
Audit Commands	357
The Audit Daemon	357
The audit Command	358
The auditreduce Command	358
The praudit Command	360
The auditconfig Command	361
Audit Files	363
The /etc/system File	363
The audit_class File	363
The audit_control File	364
The audit_data File	365
The audit_event File	365

The audit_startup Script	365
The audit_user File	366
The audit_warn Script	367
Audit Flags	368
Definitions of Audit Flags	368
Audit Flag Syntax	369
Prefixes to Modify Audit Flags	370
Audit Policies	371
Process Audit Characteristics	371
Audit Trail	372
More About the Audit Files	372
Audit Record Structure	374
Audit Token Formats	374
acl Token	376
arbitrary Token	376
arg Token	377
attr Token	378
exec_args Token	378
exec_env Token	379
exit Token	380
file Token	380
groups Token (Obsolete)	381
header Token	381
in_addr Token	382
ip Token	382
ipc Token	383
ipc_perm Token	384
iport Token	384
newgroups Token	385
opaque Token	385
path Token	386
process Token	386
return Token	388
seq Token	388
socket Token	389
subject Token	389
text Token	390
trailer Token	391

Device Allocation Reference	392
Components of the Device-Allocation Mechanism	392
Using the Device-Allocation Commands	392
The Allocate Error State	393
The <code>device_maps</code> File	393
The <code>device_allocate</code> File	394
Device-Clean Scripts	396
How the Device-Allocation Mechanism Works	398

<b>Glossary</b>	<b>401</b>
-----------------	------------

<b>Index</b>	<b>409</b>
--------------	------------



# Preface

---

*System Administration Guide: Security Services* is part of a multivolume set that covers a significant part of the Solaris™ system administration information. This book assumes that you have already installed the SunOS™ 5.9 operating system, and you have set up any networking software that you plan to use. The SunOS 5.9 operating system is part of the Solaris 9 product family, which also includes many features, such as the Solaris Common Desktop Environment (CDE).

---

**Note** – The Solaris operating environment runs on two types of hardware, or platforms—SPARC™ and IA. The Solaris operating environment runs on both 64-bit and 32-bit address spaces. The information in this document pertains to both platforms and address spaces unless called out in a special chapter, section, note, bullet, figure, table, example, or code example.

---

---

## Who Should Use This Book

This book is intended for anyone who is responsible for administering one or more systems that run the Solaris 9 release. To use this book, you should have one to two years of UNIX® system administration experience. Attending training courses in UNIX system administration might be helpful.

---

# How the System Administration Volumes Are Organized

Here is a list of the topics that are covered by the volumes of 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, removable media, managing software (packages and patches), disks and devices, file systems, and backing up and restoring data
<i>System Administration Guide: Advanced Administration</i>	Printing services, terminals and modems, system resources (disk quotas, accounting, and crontabs), system processes, and troubleshooting Solaris software problems
<i>System Administration Guide: IP Services</i>	TCP/IP networks, IPv4 and IPv6, DHCP, IP Security, Mobile IP, and IP Network Multipathing
<i>System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)</i>	DNS, NIS, and LDAP naming and directory services
<i>System Administration Guide: Naming and Directory Services (FNS and NIS+)</i>	FNS and NIS+ naming and directory services
<i>System Administration Guide: Resource Management and Network Services</i>	Resource management, remote file systems, mail, SLP, and PPP
<i>System Administration Guide: Security Services</i>	Auditing, PAM, RBAC, and SEAM

---

## Related Books

This is a list of related documentation that is referred to in this book.

- Carasik, Anne. *UNIX Secure Shell*. McGraw Hill, 1999.
- Cheswick, William R. and Steven M. Bellovin. *Firewalls and Internet Security*. Addison Wesley, 1994.

---

## Accessing Sun Documentation Online

The docs.sun.com<sup>SM</sup> Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

---

## Typographic Conventions

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

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
<b>AaBbCc123</b>	What you type, contrasted with on-screen computer output	<code>machine_name% su</code> Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type <b>rm</b> <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You must be <i>root</i> to do this.

---

## Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

**TABLE P-2** Shell Prompts

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

## Security Services (Overview)

---

This book focuses on the Solaris™ operating environment features that can help make a site more secure. This book is intended for system administrators and users of these security features. This is a list of the overview information in this chapter.

- “Introduction to Security Services” on page 21
- “Authentication” on page 22
- “Access Control” on page 22
- “Secure Communication” on page 23
- “Auditing” on page 23

---

### Introduction to Security Services

To help a site secure its computing environment, the Solaris operating environment software provides the following:

- **Authentication** – The ability to securely identify a user, requiring the user’s name and some form of proof (typically a password)
- **Access Control** – The ability to restrict users to only those parts of the system that are necessary for their job
- **Secure Communication** – The ability to ensure that authenticated parties can communicate without interception, modification, or spoofing
- **Auditing** – The ability to identify the source of security changes to the system, including file access, security-related system calls, and authentication failures

For a general discussion of system security, see Chapter 14.

---

## Authentication

Authentication is a mechanism that identifies a user or service based on predefined criteria. Authentication systems range from simple name-password pairs to more elaborate challenge-response systems, such as smart cards and biometrics. Strong authentication mechanisms rely on a user supplying information that only that person knows, such as a user name, and something that can be verified, such as a smart card or fingerprint. The Solaris operating environment features for authentication include the following:

- **Secure RPC** – An authentication technique that is based on the Diffie-Hellman method. This topic is covered in “Overview of Secure RPC” on page 27.
- **Pluggable Authentication Module (PAM)** – A framework that enables various authentication technologies to be plugged in without disturbing system entry services, such as `login` or `ftp`. See “PAM (Overview)” on page 36.
- **Sun Enterprise Authentication Module (SEAM)** – A client/server architecture that provides authentication with encryption. See Chapter 6.
- **Smart Card** – A plastic card with a microprocessor and memory that can be used with a card reader to access systems. See *Solaris Smartcard Administration Guide*.
- **Login Administration Tools** – Various commands for administering a user’s ability to log in or to abort a session. See Chapter 16.

---

## Access Control

Access control enables users or administrators to restrict the users who are permitted access to resources on the system. The Solaris operating environment features for access control include the following:

- **UNIX® permissions** – Attributes of a file or directory that control the users and groups that are permitted to read, write, or execute a file, or search a directory. See Chapter 15.
- **Role-Based Access Control (RBAC)** – An architecture for creating special, restricted user accounts that are permitted to perform specific security-related tasks. See Chapter 17.
- **Device Allocation** – A facility that enables restriction on who can use a device, such as a floppy or CD-ROM drive. The facility ensures that a device is used by only one qualified user at a time. See “Managing Device Allocation” on page 352.
- **Security Enhancement** – Through the use of scripts, many system files and parameters can be adjusted to reduce security risks. See Chapter 20.

---

## Secure Communication

The basis of secure communication is requiring authentication with encryption. Authentication helps ensure that the source and destination are the intended parties. Encryption codes the communication at the source and decodes it at the target to prevent intruders from reading any transmissions that they might manage to intercept. The Solaris operating environment features for secure communication include the following:

- **Sun™ Enterprise Authentication Module (SEAM)** – A client/server architecture that provides encryption with authentication. See Chapter 6.
- **Internet Protocol Security Architecture (IPsec)** – An architecture that provides IP datagram protection including confidentiality, strong integrity of the data, partial sequence integrity (replay protection), and data authentication. See “IPsec (Overview)” in *System Administration Guide: IP Services*.
- **Solaris Secure Shell** – A protocol for protecting data transfers and interactive user network sessions from eavesdropping, session hijacking, and man-in-the-middle attacks. Strong authentication is provided through public key cryptography. X windows services and other network services can be tunneled safely over Secure Shell connections for additional protection. See Chapter 4.

---

## Auditing

Auditing is a fundamental concept of system security and maintainability. Auditing is the process of examining the history of actions and events on a system to find out what happened. Auditing entails keeping a log of what was done, by whom, when it was done, and what was affected. For more information on Solaris operating environment auditing, see Chapter 22.





## Authentication Services Topics

---

Chapter 3	Provides information about Diffie-Hellman authentication and the Pluggable Authentication Module (PAM) framework.
Chapter 4	Provides an introduction to Solaris Secure Shell, as well as step-by-step instructions.
Chapter 5	Provides a description of the files that are used to configure Solaris Secure Shell.
Chapter 6	Provides overview information about Sun Enterprise Authentication Mechanism (SEAM).
Chapter 7	Provides a list of information that needs to be gathered and issues that need to be resolved before you configure SEAM.
Chapter 8	Provides step-by-step instructions for configuring SEAM.
Chapter 9	Provides a list of SEAM error messages, how to fix the conditions that generate the messages, and how to troubleshoot some error conditions.
Chapter 10	Provides step-by-step instructions for administering SEAM principles and policies with the <code>gkadmin</code> GUI and at the command line.
Chapter 11	Provides user instructions for SEAM.
Chapter 12	Provides reference information about SEAM.



## Using Authentication Services (Tasks)

---

The first part of this chapter provides information about the Diffie-Hellman authentication mechanism that can be used with Secure RPC. The second part covers the Pluggable Authentication Module (PAM) framework. PAM provides a method to “plug-in” authentication services and provides support for multiple authentication services.

This is a list of the step-by-step instructions in this chapter.

- “Overview of Secure RPC” on page 27
- “Administering Diffie-Hellman Authentication” on page 32
- “PAM (Overview)” on page 36
- “PAM (Tasks)” on page 38
- “PAM (Reference)” on page 42

---

### Overview of Secure RPC

Secure RPC is an authentication method that authenticates both the host and the user who is making a request for a service. Secure RPC uses the Diffie-Hellman authentication mechanism. This authentication mechanism uses DES encryption. Applications that use Secure RPC include NFS and the NIS+ name service.

### NFS Services and Secure RPC

NFS enables several hosts to share files over the network. Under the NFS service, a server holds the data and resources for several clients. The clients have access to the file systems that the server shares with the clients. Users who are logged in to the client machines can access the file systems by mounting them from the server. To the user on the client machine, it appears as if the files are local to the client. One of the

most common uses of NFS is to allow systems to be installed in offices, while keeping all user files in a central location. Some features of the NFS service, such as the `mount -nosuid` option, can be used to prohibit the opening of devices, as well as file systems, by unauthorized users.

The NFS service uses Secure RPC to authenticate users who make requests over the network. This process is known as Secure NFS. The authentication mechanism, `AUTH_DH`, uses DES encryption with Diffie-Hellman authentication to ensure authorized access. The `AUTH_DH` mechanism has also been called `AUTH_DES`.

- How to set up and administer Secure NFS in “Administering the Secure NFS System” in *System Administration Guide: Resource Management and Network Services*
- How to set up the NIS+ tables and enter names in the `cred` table in *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*
- An outline of the steps that are involved in RPC authentication in “Implementation of Diffie-Hellman Authentication” on page 29

## DES Encryption

The Data Encryption Standard (DES) encryption functions use a 56-bit key to encrypt data. If two credential users (or principals) know the same DES key, they can communicate in private, by using the key to encipher and decipher text. DES is a relatively fast encryption mechanism. A DES chip makes the encryption even faster. However, if the chip is not present, a software implementation is substituted.

The risk of using just the DES key is that an intruder can collect enough cipher-text messages that were encrypted with the same key to be able to discover the key and decipher the messages. For this reason, security systems such as Secure NFS change the keys frequently.

## Kerberos Authentication

Kerberos is an authentication system that was developed at MIT. Encryption in Kerberos is based on DES. Kerberos V4 support is no longer supplied as part of Secure RPC. However, a client-side implementation of Kerberos V5, which uses `RPCSEC_GSS`, is included with this release. For more information see Chapter 6.

## Diffie-Hellman Authentication

The Diffie-Hellman (DH) method of authenticating a user is nontrivial for an intruder to crack. The client and the server each have their own private key (sometimes called a secret key) which they use together with the public key to devise a common key. They use the common key to communicate with each other, by using an agreed-upon encryption/decryption function (such as DES). This method was identified as DES authentication in previous Solaris releases.

Authentication is based on the ability of the sending system to use the common key to encrypt the current time, which the receiving system can decrypt and check against its current time. Make sure that you synchronize the time on the client and the server.

The public keys and private keys are stored in an NIS or NIS+ database. NIS stores the keys in the `publickey` map. NIS+ stores the keys in the `cred` table. These files contain the public key and the private key for all potential users.

The system administrator is responsible for setting up NIS maps or NIS+ tables, and generating a public key and a private key for each user. The private key is stored in encrypted form with the user's password. This process makes the private key known only to the user.

## Implementation of Diffie-Hellman Authentication

This section describes the series of transactions in a client-server session that use DH authorization (`AUTH_DH`).

### *Generating the Public Keys and Secret Keys*

Sometime prior to a transaction, the administrator runs either the `newkey` or `nisaddcred` command to generate a public key and a secret key. Each user has a unique public key and secret key. The public key is stored in a public database. The secret key is stored in encrypted form in the same database. To change the key pair, use the `chkey` command.

### *Running the keylogin Command*

Normally, the login password is identical to the secure RPC password. In this case, the `keylogin` command is not required. However, if the passwords are different, the users have to log in, and then run a `keylogin` command explicitly.

The `keylogin` command prompts the user for a secure RPC password and uses the password to decrypt the secret key. The `keylogin` command then passes the decrypted secret key to a program called the `keyserver`. The `keyserver` is an RPC service with a local instance on every computer. The `keyserver` saves the decrypted secret key and waits for the user to initiate a secure RPC transaction with a server.

If both the login password and the RPC password are the same, the login process passes the secret key to the `keyserver`. If the passwords are required to be different and the user must always run the `keylogin` command, then the `keylogin` command can be included in the user's environment configuration file, such as the `~/ .login`, `~/ .cshrc`, or `~/ .profile` file. Then, the `keylogin` command runs automatically whenever the user logs in.

### *Generating the Conversation Key*

When the user initiates a transaction with a server, the following occurs:

1. The keyserver randomly generates a conversation key.
2. The kernel uses the conversation key to encrypt the client's time stamp (among other things).
3. The keyserver looks up the server's public key in the public key database (see the `publickey(4)` man page).
4. The keyserver uses the client's secret key and the server's public key to create a common key.
5. The keyserver encrypts the conversation key with the common key.

### *First Contact With the Server*

The transmission, which includes the encrypted time stamp and the encrypted conversation key, is then sent to the server. The transmission includes a credential and a verifier. The credential contains three components:

- The client's net name
- The conversation key, which is encrypted with the common key
- A "window," which is encrypted with the conversation key

The window is the difference in time that the client says should be allowed between the server's clock and the client's time stamp. If the difference between the server's clock and the time stamp is greater than the window, the server rejects the client's request. Under normal circumstances, this rejection will not happen, because the client first synchronizes with the server before starting the RPC session.

The client's verifier contains the following:

- The encrypted time stamp
- An encrypted verifier of the specified window, which is decremented by 1

The window verifier is needed in case somebody wants to impersonate a user and writes a program that, instead of filling in the encrypted fields of the credential and verifier, just stuffs in random bits. The server will decrypt the conversation key into some random key, and use it to try to decrypt the window and the time stamp. The result will be random numbers. After a few thousand trials, however, there is a good chance that the random window/time stamp pair will pass the authentication system. The window verifier makes guessing the right credential much more difficult.

### *Decrypting the Conversation Key*

When the server receives the transmission from the client, the following occurs:

1. The keyserver that is local to the server looks up the client's public key in the public key database.

2. The keyserver uses the client's public key and the server's secret key to deduce the common key, which is the same common key that is computed by the client. Only the server and the client can calculate the common key because the calculation requires knowing one of the secret keys.
3. The kernel uses the common key to decrypt the conversation key.
4. The kernel calls the keyserver to decrypt the client's time stamp with the decrypted conversation key.

### *Storing Information on the Server*

After the server decrypts the client's time stamp, it stores four items of information in a credential table:

- The client's computer name
- The conversation key
- The window
- The client's time stamp

The server stores the first three items for future use. The server stores the time stamp to protect against replays. The server accepts only time stamps that are chronologically greater than the last time stamp seen, so any replayed transactions are guaranteed to be rejected.

---

**Note** – Implicit in these procedures is the name of the caller, who must be authenticated in some manner. The keyserver cannot use DES authentication to authenticate the caller because it would create a deadlock. To solve this problem, the keyserver stores the secret keys by user ID (UID) and grants requests only to local root processes.

---

### *Returning the Verifier to the Client*

The server returns a verifier to the client, which includes the following:

- The index ID, which the server records in its credential cache
- The client's time stamp minus 1, which is encrypted by the conversation key

The reason for subtracting 1 from the time stamp is to ensure that the time stamp is invalid and that it cannot be reused as a client verifier.

### *Client Authenticates the Server*

The client receives the verifier and authenticates the server. The client knows that only the server could have sent the verifier because only the server knows what time stamp the client sent.

### *Additional Transactions*

With every transaction after the first transaction, the client returns the index ID to the server in its second transaction and sends another encrypted time stamp. The server sends back the client's time stamp minus 1, which is encrypted by the conversation key.

---

## Administering Diffie-Hellman Authentication

A system administrator can implement policies that help secure the network. The level of security that is required differs with each site. This section provides instructions for some tasks that are associated with network security.

### ▼ How to Restart the Keyserver

1. **Become superuser or assume an equivalent role.**
2. **Verify whether the `keyserv` daemon (the keyserver) is running.**

```
# ps -ef | grep keyserv
root  100      1 16  Apr 11  ?          0:00 /usr/sbin/keyserv
root  2215     2211  5 09:57:28 pts/0    0:00 grep keyserv
```

3. **Start the keyserver if it isn't running.**

```
# /usr/sbin/keyserv
```

### ▼ How to Set Up a Key in NIS+ Credentials for Diffie-Hellman Authentication

For detailed description of NIS+ security, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)*.

1. **Become superuser or assume an equivalent role.**
2. **Edit the `/etc/nsswitch.conf` file, and add the following line:**

```
publickey: nisplus
```

3. **Initialize the NIS+ client.**



```
# nisinit -cH hostname
```

*hostname* is the name of a trusted NIS+ server that contains an entry in its tables for the client machine.

4. Add the client to the `cred` table by typing the following commands:

```
# nisaddcred local
# nisaddcred des
```

5. Verify the setup by using the `keylogin` command.

If you are prompted for a password, the procedure has succeeded.

## Example—Setting Up a New Key for `root` on a NIS+ Client

The following example uses the host `pluto` to set up `earth` as an NIS+ client. You can ignore the warnings. The `keylogin` command is accepted, verifying that `earth` is correctly set up as a secure NIS+ client.

```
# nisinit -cH pluto
NIS Server/Client setup utility.
This machine is in the North.Abc.COM. directory.
Setting up NIS+ client ...
All done.
# nisaddcred local
# nisaddcred des
DES principal name : unix.earth@North.Abc.COM
Adding new key for unix.earth@North.Abc.Com (earth.North.Abc.COM.)

Network password: xxx      Press Return
Warning, password differs from login password.
Retype password: xxx      Press Return

# keylogin
Password:
#
```

## ▼ How to Set Up a New User Key That Uses NIS+ Credentials for Diffie-Hellman Authentication

1. Add the user to the `cred` table on the root master server by typing the following command:

```
# nisaddcred -p unix.UID@domain-name -P username.domain-name. des
```

Note that, in this case, the *username.domain-name* must end with a dot (`.`)

2. Verify the setup by logging in as the client and typing the `keylogin` command.

## Example—Setting Up a New Key for an NIS+ User

The following example shows DES security authorization being given to a user that is named `george`.

```
# nisaddcred -p unix.1234@North.Abc.com -P george.North.Abc.COM. des
DES principal name : unix.1234@North.Abc.COM
Adding new key for unix.1234@North.Abc.COM (george.North.Abc.COM.)

Password:
Retype password:

# rlogin rootmaster -l george
# keylogin
Password:
#
```

## ▼ How to Set Up a root Key by Using NIS Credentials With Diffie-Hellman Authentication

1. Become superuser on the client or assume an equivalent role.
2. Edit the `/etc/nsswitch.conf` file, and add the following line:

```
publickey: nis
```

3. Create a new key pair by using the `newkey` command.

```
# newkey -h hostname
hostname is the name of the client.
```

## Example—Setting Up an NIS+ Client to Use Diffie-Hellman Security

The following example shows how to set up `earth` as a secure NIS client.

```
# newkey -h earth
Adding new key for unix.earth@North.Abc.COM
New Password:
Retype password:
Please wait for the database to get updated...
Your new key has been successfully stored away.
#
```

## ▼ How to Create a New User Key That Uses NIS Credentials with Diffie-Hellman Authentication

1. Log in to the server as superuser or assume an equivalent role.

Only system administrators, when logged in to the NIS+ server, can generate a new key for a user.

## 2. Create a new key for a user.

```
# newkey -u username
```

*username* is the name of the user. The system prompts for a password. You can type a generic password. The private key is stored in an encrypted form using the generic password.

```
# newkey -u george
Adding new key for unix.12345@Abc.North.Acme.COM
New Password:
Retype password:
Please wait for the database to get updated...
Your new key has been successfully stored away.
#
```

## 3. Tell the user to log in and type the `chkey -p` command.

This command allows the user to re-encrypt his or her private key with a password known only to the user.

```
earth% chkey -p
Updating nis publickey database.
Reencrypting key for unix.12345@Abc.North.Acme.COM
Please enter the Secure-RPC password for george:
Please enter the login password for george:
Sending key change request to pluto...
#
```

---

**Note** – The `chkey` command can be used to create a new key-pair for a user.

---

## ▼ How to Share and Mount Files With Diffie-Hellman Authentication

### *Prerequisite*

The Diffie-Hellman `publickey` authentication must be enabled on the network. See “How to Set Up a Key in NIS+ Credentials for Diffie-Hellman Authentication” on page 32 and “How to Set Up a root Key by Using NIS Credentials With Diffie-Hellman Authentication” on page 34.

*To share a file system with Diffie-Hellman authentication:*

1. **Become superuser or assume an equivalent role.**
2. **Share the file system with Diffie-Hellman authentication.**

```
# share -F nfs -o sec=dh /filesystem
```

*To mount a file system with Diffie-Hellman authentication:*

1. **Become superuser or assume an equivalent role.**
2. **Mount the file system with Diffie-Hellman authentication.**

```
# mount -F nfs -o sec=dh server:resource mountpoint
```

The `-o sec=dh` option mounts the file system with `AUTH_DH` authentication.

---

## PAM (Overview)

The Pluggable Authentication Module (PAM) framework lets you “plug in” new authentication technologies without changing system entry services such as `login`, `ftp`, `telnet`, and so on. You can also use PAM to integrate UNIX login with other security mechanisms like DCE or Kerberos. Mechanisms for account, session, and password management can also be “plugged in” by using this framework.

## Benefits of Using PAM


The PAM framework allows you to choose any combination of system entry services (`ftp`, `login`, `telnet`, or `rsh`, for example) for user authentication. Some benefits that PAM provides are as follows:

- Flexible configuration policy
  - Per application authentication policy
  - The ability to choose a default authentication mechanism
  - Multiple passwords on high-security systems
- Ease of use for the end user
  - No retyping of passwords if they are the same for different mechanisms
  - The ability to use a single password for multiple authentication methods with the password-mapping feature, even if the passwords that are associated with each authentication method are different
  - The ability to prompt the user for passwords for multiple authentication methods without having the user enter multiple commands
- The ability to pass optional parameters to the user authentication services

## PAM Components

The PAM software consists of a library, several modules, and a configuration file. New versions of several commands or daemons that take advantage of the PAM interfaces are also included.

The following figure below illustrates the relationship between the applications, the PAM library, the `pam.conf` file, and the PAM modules.



**FIGURE 3-1** How PAM Works

The applications (`ftp`, `telnet`, and `login`) use the PAM library to access the appropriate module. The `pam.conf` file defines which modules to use, and in what order they are to be used with each application. Responses from the modules are passed back through the library to the application.

The following sections describe the relationship between the PAM components and the applications.

## PAM Library

The PAM library, `/usr/lib/libpam`, provides the framework to load the appropriate modules and manage the stacking process. The PAM library provides a generic structure to which all of the modules can plug in.

## Stacking Feature

The PAM framework provides a method for authenticating users with multiple services by using *stacking*. Depending on the configuration, the user can be prompted for passwords for each authentication method. The order in which the authentication services are used is determined through the PAM configuration file.

## Password-Mapping Feature

The stacking feature can require that a user remember several passwords. With the *password-mapping* feature, the primary password is used to decrypt the other passwords so that the user doesn't need to remember or enter multiple passwords. The other option is to synchronize the passwords across each authentication mechanism. Note that this strategy could increase the security risk, because the security of each mechanism is limited by the least secure password method that is used in the stack.

## Changes to PAM for the Solaris 9 Release

The Solaris 9 release includes several enhancements to the PAM service. The following list highlights the most important changes:

- To accommodate proper stacking, the `pam_unix` module is broken into single service modules. These modules provide the same capabilities as in the existing module. The capabilities are provided by the following modules: `pam_authtok_get`, `pam_authtok_check`, `pam_authtok_store`, `pam_unix_auth`, `pam_dhkeys` and `pam_passwd_auth`. See “PAM Modules” on page 42 for information about the new modules.
- New PAM services are included: `cron`, `dtsession`, `ppp` and `ssh`. See “Valid Service Names” on page 45 for information about the new services.
- The PAM configuration file was updated to include the new modules and services. See “Generic `pam.conf` File” on page 47 for information about the configuration file.

---

## PAM (Tasks)

This section discusses some tasks that might be required to make the PAM framework fully functional. In particular, you should be aware of some security issues that are associated with the PAM configuration file.

## PAM (Task Map)

Task	Description	For Instructions
Plan for your PAM Installation	Consider configuration issues and make decisions about them before you start the software configuration process.	“Planning for PAM” on page 39
Add new PAM modules	Sometimes, site-specific modules must be written and installed to cover requirements that are not part of the generic software. This procedure covers the installation process.	“How to Add a PAM Module” on page 40
Block access through <code>~/ .rhosts</code>	Steps to further increase security by preventing access through <code>~/ .rhosts</code> .	“How to Prevent Unauthorized Access From Remote Systems With PAM” on page 41
Initiate error reporting	Steps to start the reporting of PAM error messages through <code>syslog</code> .	“How to Initiate PAM Error Reporting” on page 41

### Planning for PAM

When you are deciding how best to use PAM in your environment, start by focusing on these issues:

- Determine what your needs are, especially which modules you should select.
- Identify the services that need special attention. Use `OTHER` if appropriate.
- Decide on the order in which the modules should be run.
- Select the control flag for each module.
- Choose any options that are necessary for each module.

Here are some suggestions to consider before you change the PAM configuration file:

- Use the `OTHER` entry for each module type so that every application does not have to be included.
- Make sure to consider the security implications of the `sufficient` and `optional` control flags.
- Review the man pages that are associated with the modules to understand how each module functions, what options are available, and the interactions between stacked modules.



---

**Caution** – If the PAM configuration file is misconfigured or becomes corrupted, it is possible that even superuser would be unable to log in. Since the `sulogin` command does not use PAM, superuser would then be required to boot the machine into single-user mode and fix the problem.

---

After you change the `/etc/pam.conf` file, review it as much as possible while you are still logged in as superuser. Test all the commands that might have been affected by your changes. For example, if you added a new module to the `telnet` service, use the `telnet` command and verify that the changes you made make the service behave as expected.

## ▼ How to Add a PAM Module

1. **Become superuser or assume an equivalent role.**
2. **Determine which control flags and other options should be used.**  
Refer to “PAM Modules” on page 42 information on the modules.
3. **Copy the new module to `/usr/lib/security/sparcv9`.**  
In the Solaris 8 release, the module should be copied to `/usr/lib/security`.
4. **Set the permissions so that the module file is owned by `root` and that permissions are `555`.**
5. **Edit the PAM configuration file, `/etc/pam.conf`, and add this module to the appropriate services.**

## Verification

It is very important to do some testing *before* the system is rebooted in case the configuration file is misconfigured. Run `rlogin`, `su`, and `telnet` before you reboot the system. If the service is a daemon that is spawned only once when the system is booted, it might be necessary to reboot the system before you can verify that the module has been added.



## ▼ How to Prevent Unauthorized Access From Remote Systems With PAM

Remove the `rlogin auth rhosts_auth.so.1` entry from the PAM configuration file. This step prevents the reading of the `~/.rhosts` files during an `rlogin` session and therefore prevents unauthenticated access to the local system from remote systems. All `rlogin` access requires a password, regardless of the presence or contents of any `~/.rhosts` or `/etc/hosts.equiv` files.

---

**Note** – To prevent other unauthenticated access to the `~/.rhosts` files, remember to disable the `rsh` service. The best way to disable a service is to remove the service entry from the `/etc/inetd.conf` file. Changing the PAM configuration file does not prevent the service from being started.

---

## ▼ How to Initiate PAM Error Reporting

1. **Edit the `/etc/syslog.conf` file to add any of the following entries for PAM error reporting:**
  - `auth.alert` - Messages about conditions that should be fixed immediately
  - `auth.crit` - Critical messages
  - `auth.err` - Error messages
  - `auth.info` - Informational messages
  - `auth.debug` - Debugging messages
2. **Restart the `syslog` daemon, or send a `SIGHUP` signal to it to activate the PAM error reporting.**

### Example—Initiating PAM Error Reporting

In the following example, all alert messages are displayed on the console. Critical messages are mailed to root. Informational and debug messages are added to the `/var/log/pamlog` file.

```
auth.alert    /dev/console
auth.crit     'root'
auth.info;auth.debug    /var/log/pamlog
```

Each line in the log contains a time stamp, the name of the system that generated the message, and the message itself. The `pamlog` file is capable of logging a large amount of information.

---

## PAM (Reference)

PAM uses run-time pluggable modules to provide authentication for system entry services. These modules are broken into four different types, based on their function:

- authentication
- account management
- session management
- password management

A stacking feature is provided to let you authenticate users through multiple services. Also provided is a password-mapping feature to not require that users remember multiple passwords.

## PAM Modules

Each PAM module implements a specific mechanism. When you set up PAM authentication, you need to specify both the module and the module type, which defines what the module does. More than one module type (auth, account, session, or password) can be associated with each module.

The following table describes each PAM module, and includes the module name and the module file name. The path of each module is determined by the instruction set that is available in the Solaris release that is installed. The default path to the modules is `/usr/lib/security/$ISA`. The value for `$ISA` could be `sparc` or `i386`. See the `isalist(5)` man page for more information.

**TABLE 3-1** PAM Modules

Module Name and Module File Name	Description
authtok_check pam_authtok_check.so.1	Provides support for password management. This module performs various checks on passwords, such as checking the length of the password, for circular shift of the login name, for password complexity, and for the amount of variation between new passwords and old passwords. See <code>pam_authtok_check(5)</code> for more information.
authtok_get pam_authtok_get.so.1	Provides password prompting for authentication and password management. See <code>pam_authtok_get(5)</code> for more information.
authtok_store pam_authtok_store.so.1	Provides support for authentication only. This module updates the authentication token for the user. After the successful update, the module stores the token in the specified repository or default repository. See <code>pam_authtok_store(5)</code> for more information.

**TABLE 3-1** PAM Modules (Continued)

Module Name and Module File Name	Description
dhkeys pam_dhkeys.so.1	Provides support for Diffie-Hellman key management in authentication. This module supports Secure RPC authentication and Secure RPC authentication token management. See <code>pam_dhkeys(5)</code> for more information.
dial_auth pam_dial_auth.so.1	Can only be used for authentication. This module uses that is data stored in the <code>/etc/dialups</code> and <code>/etc/d_passwd</code> files for authentication. This module is mainly used by the <code>login</code> command. See <code>pam_dial_auth(5)</code> for more information.
krb5 pam_krb5_auth.so.1	Provides support for authentication, account management, session management, and password management. Kerberos credentials are used for authentication. See <code>pam_krb5(5)</code> for more information.
ldap pam_ldap.so.1	Provides support for authentication and password management. Data from an LDAP server are used for authentication. See <code>pam_ldap(5)</code> for more information.
projects pam_projects.so.1	Provides support for account management. See <code>pam_projects(5)</code> for more information.
rhosts_auth pam_rhosts_auth.so.1	Can only be used for authentication. This module uses data that is stored in the <code>~/.rhosts</code> and <code>/etc/host.equiv</code> files through the <code>ruserok()</code> routine. This module is mainly used by the <code>rlogin</code> and <code>rsh</code> commands. See <code>pam_rhosts_auth(5)</code> for more information.
roles pam_roles.so.1	Provides support for account management only. The RBAC <code>user_attr</code> database determines the roles a user can assume. See <code>pam_roles(5)</code> for more information.
sample pam_sample.so.1	Provides support for authentication, account management, session management, and password management. Used for testing. See <code>pam_sample(5)</code> for more information.
smartcard pam_smartcard.so.1	Provides support for authentication only. See <code>pam_smartcard(5)</code> for more information.
unix pam_unix.so.1	Provides support for authentication, account management, session management, and password management. Any of the four module type definitions can be used with this module. This module uses UNIX passwords for authentication.  In the Solaris environment, the selection of appropriate name services to get password records is controlled through the <code>/etc/nsswitch.conf</code> file. See <code>pam_unix(5)</code> for more information.

**TABLE 3-1** PAM Modules (Continued)

Module Name and Module File Name	Description
unix_account pam_unix_account.so.1	Provides support for account management. This module retrieves password aging information from the repository specified in the <code>nsswitch.conf</code> file and verifies that the password and the user's account have not expired. See <code>pam_unix_account(5)</code> for more information.
unix_auth pam_unix_auth.so.1	Provides support for authentication. This module verifies that the password contained in the PAM handle is the correct password for the user's password in the specified repository or default repository. See <code>pam_unix_auth(5)</code> for more information.
unix_session pam_unix_session.so.1	Provides support for session management. This module initiates session management by updating the <code>/var/adm/lastlog</code> file. See <code>pam_unix_session(5)</code> for more information.

For security reasons, these module files must be owned by `root` and must not be writable through `group` or `other` permissions. If the file is not owned by `root`, PAM will not load the module.

## PAM Module Types

It is important to understand the PAM module types because they define the interface to the module. Here are the four types of run-time PAM modules:

- The *authentication modules* provide authentication for the users and allow for credentials to be set, refreshed, or destroyed. They provide a valuable administration tool for user identification.
- The *account modules* check for password aging, account expiration, and access hour restrictions. After the user is identified through the authentication modules, the account modules determine if the user should be given access.
- The *session modules* manage the opening and the closing of an authentication session. They can log activity or provide for clean-up after the session is over.
- The *password modules* allow for changes to the actual password.

## PAM Configuration File

The PAM configuration file, `/etc/pam.conf`, determines the authentication services to be used, and in what order they are used. This file can be edited to select authentication mechanisms for each system entry application.

## PAM Configuration File Syntax

The PAM configuration file consists of entries with the following syntax:

*service\_name module\_type control\_flag module\_path module\_options*

<i>service_name</i>	Is the name of the service (for example, <code>ftp</code> , <code>login</code> , <code>telnet</code> ).
<i>module_type</i>	Is the module type for the service. For more information see “PAM Module Types” on page 44.
<i>control_flag</i>	Determines the continuation or failure behavior for the module.
<i>module_path</i>	Specifies the path to the library object that implements the service.
<i>module_options</i>	Specifies the options that are passed to the service modules.

You can add comments to the `pam.conf` file by starting the line with a # (pound sign). Use white spaces or tabs to delimit the fields.

---

**Note** – An entry in the PAM configuration file is ignored if one of the following conditions exist: the line has less than four fields, an invalid value is given for *module\_type* or *control\_flag*, or the named module does not exist.

---

## Valid Service Names

The following table lists some valid service names, the module types that can be used with that service, and the daemon or command that is associated with the service name.

Not all module types are appropriate for each service. For example, the `password` module type is appropriate for only the `passwd` command. Also, since the `passwd` command is not concerned with authentication, there is no `auth` module type associated with it.

**TABLE 3-2** Valid Service Names for the `/etc/pam.conf` File

Service Name	Daemon or Command	Applicable Module Types
<code>cron</code>	<code>/usr/sbin/cron</code>	<code>auth</code> , <code>account</code>
<code>dtlogin</code>	<code>/usr/dt/bin/dtlogin</code>	<code>auth</code> , <code>account</code> , <code>session</code>

**TABLE 3-2** Valid Service Names for the `/etc/pam.conf` File (Continued)

Service Name	Daemon or Command	Applicable Module Types
dtssession	<code>/usr/dt/bin/dtssession</code>	auth
ftp	<code>/usr/sbin/in.ftpd</code>	auth, account, session
init	<code>/usr/sbin/init</code>	session
login	<code>/usr/bin/login</code>	auth, account, session
passwd	<code>/usr/bin/passwd</code>	password
ppp	<code>/usr/bin/ppp</code>	auth, account, session
rexcd	<code>/usr/sbin/rpc.rexd</code>	account, session
rlogin	<code>/usr/sbin/in.rlogind</code>	auth, account, session
rsh	<code>/usr/sbin/in.rshd</code>	auth, account, session
sac	<code>/usr/lib/saf/sac</code>	session
ssh	<code>/usr/bin/ssh</code>	auth, account, session
su	<code>/usr/bin/su</code>	auth, account
telnet	<code>/usr/sbin/in.telnetd</code>	auth, account, session
ttymon	<code>/usr/lib/saf/ttymon</code>	session
uucp	<code>/usr/sbin/in.uucpd</code>	auth, account, session

## Control Flags

To determine the continuation or failure behavior from a module during the authentication process, you must select one of four *control flags* for each entry in the PAM configuration file, `/etc/pam.conf`. The control flags indicate how a successful attempt or a failed attempt through each module is handled. Even though these flags apply to all module types, the following explanation assumes that these flags are being used for authentication modules. The control flags are as follows:

- **required** - With this control flag, the module must return success in order to have the overall result be successful.

If all modules are flagged as **required**, then authentication through all modules must succeed for the user to be authenticated.

If some modules fail, then an error value from the first failed module is reported.

If a failure occurs for a module that is flagged as **required**, all modules in the stack are still tried, but failure is returned.

If none of the modules are flagged as **required**, then at least one entry for that service must succeed for the user to be authenticated.

- `requisite` - With this control flag, the module must return success for additional authentication to occur.

If a failure occurs for a module that is flagged as `requisite`, an error is immediately returned to the application, and no additional authentication is done. If the stack does not include prior modules flagged as `required` that failed, then the error from this module is returned. If an earlier module flagged as `required` has failed, the error message from the `required` module is returned.

- `optional` - If a module with this control flag fails, the overall result can be successful if another module in this stack returns success.

The `optional` control flag should be used when successful authentication in the stack is enough for a user to be authenticated. This flag should only be used if it is not important for this particular mechanism to succeed.

If your users need to have permission associated with a specific mechanism to get their work done, then you should not label it as `optional`.

- `sufficient` - If a module with this control flag is successful, skip the remaining modules in the stack, even if they are flagged as `required`.

The `sufficient` control flag indicates that one successful authentication will be enough for the user to be granted access.

More information about these control flags is provided in the following section, which describes the default `/etc/pam.conf` file.

## Generic `pam.conf` File

The generic `/etc/pam.conf` file specifies the following actions:

1. When the `login` command is run, authentication must succeed for the `pam_authok_get`, `pam_dhkeys`, `pam_auth_unix` and the `pam_dial_auth` modules.
2. For the `rlogin` command, authentication through the `pam_authok_get`, `pam_dhkeys`, and `pam_auth_unix` modules must succeed, if authentication through `pam_rhost_auth` fails.
3. The `sufficient` control flag indicates that for the `rlogin` command the successful authentication that is provided by the `pam_rhost_auth` module is sufficient, and the next entry will be ignored.
4. Most of the other commands that require authentication require successful authentication through the `pam_authok_get`, `pam_dhkeys`, and `pam_auth_unix` modules.
5. For the `rsh` command, authentication through the `pam_rhost_auth` module is flagged as `sufficient`. No other authentication is required if authentication through the `pam_rhost_auth` module succeeds.

The `OTHER` service name allows a default to be set for any other commands that require authentication that are not included in the file. The `OTHER` option makes it easier to administer the file, since many commands that are using the same module

can be covered by using only one entry. Also, the OTHER service name, when used as a “catch-all,” can ensure that each access is covered by one module. By convention, the OTHER entry is included at the bottom of the section for each module type.

Normally, the entry for the *module\_path* is “root-relative.” If the file name you enter for *module\_path* does not begin with a slash (/), the path /usr/lib/security/\$ISA is prepended to the file name. A full path name must be used for modules that are located in other directories. The values for the *module\_options* can be found in the man pages for the module. For example, the UNIX module is covered in the pam\_unix(5) man page.

```
login  auth required          pam_authtok_get.so.1
login  auth required          pam_dhkeys.so.1
login  auth required          pam_unix_auth.so.1
login  auth required          pam_dial_auth.so.1
```

In this example, the login service specifies authentication through all four authentication modules. A login command fails if any of the modules return an error.



## Using Secure Shell (Tasks)

---

This chapter covers the following topics:

- “Using Secure Shell (Task Map)” on page 51
- “Introduction to Secure Shell” on page 49
- “Using Secure Shell” on page 52

---

### Introduction to Secure Shell

Secure Shell allows a user to securely access a remote host over an unsecured network. Authentication is provided by the use of passwords, public keys, or both. All network traffic is encrypted. Thus, Secure Shell prevents a would-be intruder from being able to read an intercepted communication or from spoofing the system.

Secure Shell provides commands for remote login and remote file transfer. Secure Shell can also be used as an on-demand virtual private network (VPN) to forward X Window system traffic or individual port numbers between the local machines and remote machines over the encrypted network link.

With Secure Shell, you can perform these actions:

- Log in to another host securely over an unsecured network
- Copy files securely between the two hosts
- Run commands securely on the remote host

Solaris Secure Shell supports the two versions of the Secure Shell protocol. Version 1 is the original version of the protocol. Version 2 is more secure, and it amends some of the basic security design flaws of Version 1. As a result, Version 1 is deprecated and is provided only to assist users who are migrating to Version 2. Users are strongly discouraged from using Version 1.

---

**Note** – Hereafter in this text, v1 is used to represent Version 1, and v2 is used to represent Version 2.

---

The requirements for Secure Shell authentication are as follows:

- **User authentication** – A user can be authenticated through either of the following:
  - **Passwords** – The user supplies the account password as in the login process.
  - **Public keys** – The user can create a public/private key pair that is stored on the local host. The remote hosts are provided with the public key, which is required to complete the authentication.

The source host maintains the private key, and target hosts are provided with the public key that is needed to complete authentication. Public key authentication is a stronger authentication mechanism than password authentication, because the private key never travels over the network. The public/private key pair is stored in the user's home directory under the `.ssh` subdirectory. The following table shows the default names for the identity files, which store the public keys and private keys.

**TABLE 4-1** Naming Conventions for Identity Files

Private Key, Public Key	Cipher and Protocol Version
identity, identity.pub	RSA v1
id_rsa, id_rsa.pub	RSA v2
id_dsa, id_dsa.pub	DSA v2

- **Host authentication** – Host authentication requires the remote host to have access to the local host's public key. A copy of the local host's public key is stored in `$HOME/.ssh/known_hosts` on the remote host.

The following table shows the authentication methods, the compatible protocol versions, the local host and remote host requirements, and the relative security. Note that the default method is password-based authentication.

**TABLE 4-2** Authentication Methods for Secure Shell

Authentication Method (Protocol Version)	Local Host Requirements	Remote Host Requirements	Security Level
Password-based (v1 or v2)	user account	user account	Medium

**TABLE 4-2** Authentication Methods for Secure Shell (Continued)

Authentication Method (Protocol Version)	Local Host Requirements	Remote Host Requirements	Security Level
RSA/DSA public key (v2)	user account  private key in <code>\$HOME/.ssh/id_rsa</code> or <code>\$HOME/.ssh/id_dsa</code>  public key in <code>\$HOME/.ssh/id_rsa.pub</code> or <code>\$HOME/.ssh/id_dsa.pub</code>	user account  user's public key ( <code>id_rsa.pub</code> or <code>id_dsa.pub</code> ) in <code>\$HOME/.ssh/authorized_keys</code>	Strong
RSA public key (v1)	user account  private key in <code>\$HOME/.ssh/identity</code>  public key in <code>\$HOME/.ssh/identity.pub</code>	user account  user's public key ( <code>identity.pub</code> ) in <code>\$HOME/.ssh/authorized_keys</code>	Strong
.rhosts with RSA (v1)	user account	user account  local host name in <code>/etc/hosts.equiv</code> , <code>/etc/shosts.equiv</code> , <code>\$HOME/.rhosts</code> , or <code>\$HOME/.shosts</code>  local host public key in <code>\$HOME/.ssh/known_hosts</code> or <code>/etc/ssh/ssh_known_hosts</code>	Medium
.rhosts only (v1 or v2)	user account	user account  local host name in <code>/etc/hosts.equiv</code> , <code>/etc/shosts.equiv</code> , <code>\$HOME/.rhosts</code> , or <code>\$HOME/.shosts</code>	Weak

## Using Secure Shell (Task Map)

Task	Description	For Instructions
Create a public/private key pair	Using public/private key pairs is the preferred method for authenticating yourself and encrypting your communications.	"How to Create a Public/Private Key Pair" on page 52

Task	Description	For Instructions
Log in with Secure Shell	Encrypted Secure Shell communication is enabled by logging in remotely through a process similar to using <code>rsync</code> .	"How to Log In to Another Host with Secure Shell" on page 53
Log in without a password with Secure Shell	You can log in using Secure Shell without having to provide a password by using <code>ssh-agent</code> . The <code>ssh-agent</code> command can be run manually or from a startup script.	"How to Log in With No Password While Using <code>ssh-agent</code> " on page 54 "How to Set <code>ssh-agent</code> to Run Automatically" on page 55
Port forwarding in Secure Shell	You can specify a local port or a remote port to be used in a Secure Shell connection.	"How to Use Secure Shell Port Forwarding" on page 56
Copy files with Secure Shell	You can copy remote files securely.	"How to Copy Files With Secure Shell" on page 57
Transfer files with Secure Shell	You can log into a remote host with Secure Shell by using transfer commands similar to <code>ftp</code> .	"Transferring Files Remotely Using <code>sftp</code> " on page 58
Connect from a host inside a firewall to a host on the outside	Secure shell provides commands compatible with HTTP or SOCKS5 that can be specified in a configuration file or on the command line.	"How to Set Up Default Connections to Hosts Outside a Firewall" on page 59 "Example — Connecting to Hosts Outside a Firewall From the Command Line" on page 60

## Using Secure Shell

### ▼ How to Create a Public/Private Key Pair

The standard procedure for creating a Secure Shell public/private key pair follows. For information on additional options, see `ssh-keygen(1)`.

#### 1. Start the key generation program.

```
myLocalHost% ssh-keygen
Generating public/private rsa key pair.
Enter file in which to save the key(/home/johndoe/.ssh/id_rsa):
```

#### 2. Enter the path to the file that will hold the key.

By default, the file name `id_rsa`, which represents an RSA v2 key, appears in parentheses. You can select this file by pressing Return. Or, you can type an alternative filename.

Enter file in which to save the key(/home/johndoe/.ssh/id\_rsa): **<Return>**  
The public key name is created automatically and the string .pub is appended to the private key name.

### 3. Enter a passphrase for using your key.

This passphrase is used for encrypting your private key. A good passphrase is 10–30 characters long, mixes alphabetic and numeric characters, and avoids simple English prose and English names. A null entry means no passphrase is used, but this entry is *strongly discouraged* for user accounts. Note that the passphrase is not displayed when you type it in.

Enter passphrase(empty for no passphrase): **<Type the passphrase>**

### 4. Re-enter the passphrase to confirm it.

Enter same passphrase again: **<Type the passphrase>**  
Your identification has been saved in /home/johndoe/.ssh/id\_rsa.  
Your public key has been saved in /home/johndoe/.ssh/id\_rsa.pub.  
The key fingerprint is:  
0e:fb:3d:57:71:73:bf:58:b8:eb:f3:a3:aa:df:e0:d1 johndoe@myLocalHost

### 5. Check the results.

The key fingerprint (a colon-separated series of 2 digit hexadecimal values) is displayed. Check that the path to the key is correct. In the example, the path is /home/johndoe/.ssh/id\_rsa.pub. At this point, you have created a public/private key pair.

### 6. Copy the public key and append the key to the \$HOME/.ssh/authorized\_keys file in your home directory on the remote host.

## ▼ How to Log In to Another Host with Secure Shell

### 1. Use the ssh command, specifying the name of the remote host.

```
myLocalHost% ssh myRemoteHost
```

The first time you run ssh, a prompt questions the authenticity of the remote host, as follows.

```
The authenticity of host 'myRemoteHost' can't be established.  
RSA key fingerprint in md5 is: 04:9f:bd:fc:3d:3e:d2:e7:49:fd:6e:18:4f:9c:26  
Are you sure you want to continue connecting(yes/no)?
```

This prompt is normal, and you should type **yes** and continue. If you have used ssh in the past on this remote host and this warning message still appears, then there might be a breach in your security.

### 2. Enter the Secure Shell passphrase and the account password when you are prompted for them.

```
Enter passphrase for key '/home/johndoe/.ssh/id_rsa': <Return>  
johndoe@myRemoteHost's password: <Return>
```

```
Last login: Fri Jul 20 14:24:10 2001 from myLocalHost
myRemoteHost%
```

Conduct any transactions on the remote host. Any commands that you send and any responses that are received are encrypted.

---

**Note** – If you wish to subsequently change your passphrase, use `ssh-keygen` with the `-p` option.

---

**3. When you are finished with your remote session, type `exit` or use your usual method for exiting your shell.**

```
myRemoteHost% exit
myRemoteHost% logout
Connection to myRemoteHost closed
myLocalHost%
```

## ▼ How to Log in With No Password While Using `ssh-agent`

If you want to omit passphrase and password entry when you are using Secure Shell, you can use the agent daemon. Use the `ssh-agent` command at the beginning of the session. Then, store your private keys with the agent by using `ssh-add`. If you have different accounts on different hosts, add those keys that you intend to use in the session. You can start the agent manually when needed as described in the following procedure. Or, you can set the agent to run automatically at the start of every session as described in “How to Set `ssh-agent` to Run Automatically” on page 55.

**1. Start the agent daemon.**

The `ssh-agent` command starts the agent daemon and displays its process ID.

```
myLocalHost% eval `ssh-agent`
Agent pid 9892
myLocalHost%
```

**2. Add your private key to the agent daemon.**

The `ssh-add` command adds your private key to the agent daemon so that subsequent Secure Shell activity will not prompt you for the passphrase.

```
myLocalHost% ssh-add
Enter passphrase for /home/johndoe/.ssh/id_rsa:
Identity added: /home/johndoe/.ssh/id_rsa (/home/johndoe/.ssh/id_rsa)
myLocalHost%
```

**3. Start a Secure Shell session.**

```
myLocalHost% ssh myRemoteHost
```

## Example—Using ssh-add Options

You can use `ssh-add` to add other keys to the daemon as well. For example, you might concurrently have DSA v2, RSA v2, and RSA v1 keys. To list all keys that are stored in the daemon, use the `-l` option. To delete a single key from the daemon, use the `-d` option. To delete all keys, use the `-D` option.

```
myLocalHost% eval `ssh-agent`
Agent pid 3347
myLocalHost% ssh-add
Enter passphrase for /home/johndoe/.ssh/id_rsa:
Identity added: /home/johndoe/.ssh/id_rsa(/home/johndoe/.ssh/id_rsa)
myLocalHost% ssh-add /home/johndoe/.ssh/id_dsa
Enter passphrase for /home/johndoe/.ssh/id_dsa: <type passphrase>
Identity added:
/home/johndoe/.ssh/id_dsa(/home/johndoe/.ssh/id_dsa)
myLocalHost% ssh-add -l
md5 1024 0e:fb:3d:53:71:77:bf:57:b8:eb:f7:a7:aa:df:e0:d1
/home/johndoe/.ssh/id_rsa (RSA)
md5 1024 c1:d3:21:5e:40:60:c5:73:d8:87:09:3a:fa:5f:32:53
/home/johndoe/.ssh/id_dsa (DSA)
myLocalHost% ssh-add -d
Identity removed:
/home/johndoe/.ssh/id_rsa (/home/johndoe/.ssh/id_rsa.pub)
/home/johndoe/.ssh/id_dsa (DSA)
```

## ▼ How to Set ssh-agent to Run Automatically

If you want to avoid providing your passphrase and password whenever you use Secure Shell, you can start an agent daemon (`ssh-agent`) by using the `.dtprofile` script.

1. **To start the agent daemon automatically, add the following lines to the end of the `$HOME/.dtprofile` script:**

```
if [ "$SSH_AUTH_SOCK" = "" -a -x /usr/bin/ssh-agent ]; then
    eval `usr/bin/ssh-agent`
fi
```

2. **To terminate the Secure Shell agent daemon when you exit the CDE session, add the following to the `$HOME/.dt/sessions/sessionexit` script:**

```
if [ "$SSH_AGENT_PID" != "" -a -x /usr/bin/ssh-agent ]; then
    /usr/bin/ssh-agent -k
fi
```

This entry ensures that no one can use the Secure Shell agent after the CDE session is terminated.

3. **Start a Secure Shell session.**

```
myLocalHost% ssh myRemoteHost
```

There is no prompt for a passphrase.

## ▼ How to Use Secure Shell Port Forwarding

You can specify a local port to be forwarded to a remote host. Effectively, a socket is allocated to listen to the port on the local side. The connection from this port is made over a secure channel to the remote host. For example, you might specify port 143 to obtain email remotely with IMAP4. Similarly, a port can be specified on the remote side.

1. **To set a local port to be forwarded, specify the local port to listen to, and the remote host and port to forward to.**

```
myLocalHost% ssh -L localPort:remoteHost:remotePort
```

2. **To set a remote port to receive a secure connection, specify the remote port to listen to, and the local host and port to forward to.**

```
myLocalHost% ssh -R remotePort:localHost:localPort
```

### Example — Using Local Port Forwarding to Receive Mail

The following example demonstrates how you can use local port forwarding to receive mail securely from a remote server.

```
myLocalHost% ssh -L 9143:myRemoteHost:143 myRemoteHost
```

This command forwards connections to port 9143 on `myLocalHost` to port 143 (the IMAP v2 server port) on `myRemoteHost`. When the user launches a mail application, he or she needs to specify the local port number. An example that uses the `dtmail` command is shown in Figure 4-1.

Note that the term *localhost* in this case and in “Example — Using Remote Port Forwarding to Communicate Outside of a Firewall” on page 57 refers to the keyword that designates the user’s local host. *localhost* should not be confused with *myLocalHost*, the hypothetical host name that is used to identify a local host in the examples in this chapter.



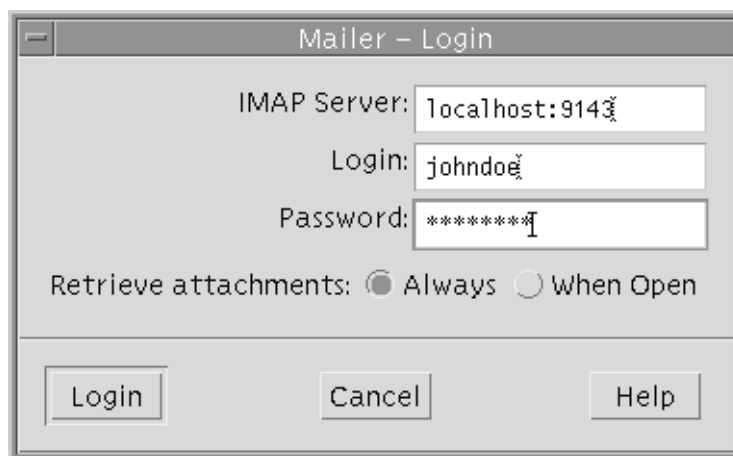


FIGURE 4-1 Specifying Port Forwarding for Email

## Example — Using Remote Port Forwarding to Communicate Outside of a Firewall

This example demonstrates how a user in an enterprise environment can forward connections from a host on an external network to a host inside a corporate firewall.

```
myLocalHost% ssh -R 9022:myLocalHost:22 myOutsideHost
```

This command forwards connections to port 9022 on myOutsideHost to port 22 (the sshd server) on the local host.

```
myOutsideHost% ssh -p 9022 localhost
myLocalHost%
```

This command demonstrates how after the remote forwarding connection has been established, the user can use ssh to connect securely from the remote host.

## ▼ How to Copy Files With Secure Shell

Use the `scp` command to copy encrypted files between hosts. You can copy encrypted files between either a local and remote host, or between two remote hosts. The command operates similarly to the `rscp` command except that it prompts for passwords. See `scp(1)` for more information.

### 1. Start the secure copy program.

Specify the source file, user name at remote destination, and destination directory.

```
myLocalHost% scp myfile.1 johndoe@myRemoteHost:~
```

## 2. Type the Secure Shell passphrase when prompted.

```
Enter passphrase for key '/home/johndoe/.ssh/id_rsa': <Return>
myfile.1      25% |*****          |      640 KB  0:20 ETA
myfile.1
```

After you type the passphrase, the progress meter is displayed. See the second line in the preceding output. The progress meter displays:

- the file name
- the percentage of the file that has been transferred at this point
- a series of asterisks that are analogous to the percentage transferred
- the quantity of data transferred
- the estimated time of arrival (ETA) of the complete file (that is, the remaining amount of time)

## Transferring Files Remotely Using `sftp`

The `sftp` command works similarly to `ftp` but uses a different set of commands. The following table lists some representative commands.

**TABLE 4-3** Interactive `sftp` Subcommands

Category	Subcommands	Description
Navigation	<code>cd path</code> ,	Changes the remote directory to <i>path</i> .
	<code>lcd path</code>	Changes the local directory to <i>path</i> .
Ownership	<code>chgrp group file</code>	Changes the group for <i>file</i> to <i>group</i> , a numeric GID.
	<code>chmod mode file</code>	Changes the permissions of <i>file</i> .
File copying	<code>get remote_file</code> [ <i>local-path</i> ]	Retrieves <i>remote_file</i> and stores it on the local host.
	<code>put local_file</code> [ <i>remote-path</i> ]	Stores a local file on the remote host.
	<code>rename old_file new_file</code>	Renames a local file.
Directory listing	<code>ls [path]</code>	Lists the contents of the remote directory.
Directory creation	<code>mkdir path</code>	Creates a remote directory.
Miscellaneous	<code>exit</code> , <code>quit</code>	Quits the <code>sftp</code> command.

## ▼ How to Set Up Default Connections to Hosts Outside a Firewall

You can use Secure Shell to make a connection from a host inside a firewall to a host on the other side of the firewall. This task is done by specifying a proxy command for `ssh` either in a configuration file or as an option on the command line. For more information see “Example — Connecting to Hosts Outside a Firewall From the Command Line” on page 60.

In general, you can customize your `ssh` interactions through a configuration file, either your own personal file `$HOME/.ssh/config` or an administrative configuration file in `/etc/ssh/ssh_config`. See `ssh_config(4)`. There are two types of proxy commands. One proxy command is for HTTP connections. The other proxy command is for SOCKS5 connections.

### 1. Specify the proxy commands and hosts in a configuration file.

Use the following syntax to add as many lines as you need:

```
[Host outside_host]  
ProxyCommand proxy_command [-h proxy_server] \  
[-p proxy_port] outside_host | %h outside_port | %p
```

Use the `Host outside_host` option to limit this proxy command specification to instances when this host (or hosts if a wildcard is used) is specified on the command line.

The designation `proxy_command` can be replaced by either of the following:

- `/usr/lib/ssh/ssh-http-proxy-connect` for HTTP connections
- `/usr/lib/ssh/ssh-socks5-proxy-connect` for SOCKS5 connections

The `-h proxy_server` and `-p proxy_port` options specify a proxy server and a proxy port, respectively. If present, they override any environment variables that specify proxy servers and ports, such as `HTTPPROXY`, `HTTPPROXYPORT`, `http_proxy` (for specifying a URL), `SOCKS5_SERVER`, and `SOCKS5_PORT`. If the options are not used, then the relevant environment variables must be set. See `ssh-socks5-proxy-connect(1)` and `ssh-http-proxy-connect(1)`.

Use `outside_host` to designate a specific host to connect to or use `%h` to specify the host on the command line. Use `outside_port` or `%p` to specify the port. Specifying `%h` and `%p` without using the `Host outside_host` option has the effect of applying the proxy command to the host argument whenever `ssh` is invoked.

### 2. Run Secure Shell, specifying the outside host.

For example, type:

```
myLocalHost% ssh myOutsideHost
```

This command looks for a proxy command specification for `myOutsideHost` in your personal configuration file. If not found, then the command looks in the system-wide configuration file, `ssh_config`. The proxy command is substituted for `ssh`.

## Example — Connecting to Hosts Outside a Firewall From the Command Line

The `-o` option for `ssh` lets you type any line that is permitted in an `ssh` configuration file. In this case the proxy command specification from the previous task is used.

Specify the proxy commands and hosts in a configuration file.

Run `ssh` and include a proxy command specification as a `-o` option. For example, type:

```
% ssh -o'Proxycommand=/usr/lib/ssh/ssh-http-proxy-connect \  
-h myProxyServer -p 8080 myOutsideHost 22' myOutsideHost
```

This command substitutes the HTTP proxy command for `ssh`, uses port 8080 and `myProxyServer` as the proxy server, and connects to port 22 on `myOutsideHost`.

## Secure Shell Administration (Reference)

---

This chapter describes how Secure Shell works from the administrator's point of view and how it is configured. This is a list of the reference information in this chapter.

- "A Typical Secure Shell Session" on page 61
- "Configuring the Secure Shell" on page 63
- "Maintaining Known Hosts on a Site-Wide Basis" on page 67
- "Secure Shell Files" on page 67

---

### A Typical Secure Shell Session

The Secure Shell daemon (`sshd`) is normally started at boot from the `/etc/init.d/sshd` script. The daemon listens for connections from clients. A Secure Shell session begins when the user runs the `ssh`, `scp`, or `sftp` command. A new `sshd` daemon is forked for each incoming connection. The forked daemons handle key exchange, encryption, authentication, command execution, and data exchange with the client. These session characteristics are determined by client-side configuration files and server configuration files, and potentially command-line parameters. The client and server must authenticate themselves to each other. After successful authentication, the user can execute commands remotely and copy data between hosts.

### Session Characteristics

The server-side behavior of the `sshd` daemon is controlled by keyword settings in the `/etc/ssh/sshd_config` file and potentially the command-line options when `sshd` is started. For example, `sshd_config` controls which types of authentication are permitted for accessing the server.

The behavior on the client side is controlled by the Secure Shell parameters in this order of precedence:

- Command-line options
- User's configuration file (`$HOME/.ssh/config`)
- System-wide configuration file (`/etc/ssh/ssh_config`)

For example, a user can override a system-wide configuration `Cipher` that is set to `blowfish` by specifying `-c 3des` on the command line.

## Authentication

The steps in the authentication process for Secure Shell are as follows:

1. The user runs the `ssh`, `scp`, or `sftp` command.
2. The client and server agree on a shared session key.

In v1, the remote host sends its host (RSA) key and a server (RSA) key to the client. Note that the server key is typically generated every hour and stored in memory only. The client checks that the remote host key is stored in the `$HOME/.ssh/known_hosts` file on the local host. The client then generates a 256 bit random number and encrypts it with the remote host's host key and server key. The encrypted random number is used as a session key to encrypt all further communications in the session.

In v2, the remote host uses DSA in its host key and does not generate a server key. Instead, the shared session key is derived through a Diffie-Hellman agreement.
3. The local and remote hosts authenticate each other.

In v1, the client can use `.rhosts`, `.rhosts` with RSA, RSA challenge-response, or password-based authentication. In v2, only `.rhosts`, DSA, and password-based authentication are permitted.

## Command Execution and Data Forwarding

After authentication is complete, the user can use Secure Shell, generally by requesting a shell or executing a command. Through the `ssh` options, the user can make requests, such as allocating a pseudo-tty, forwarding X11 connections or TCP/IP connections, or enabling an `ssh-agent` over a secure connection. The basic components of a user session are as follows:

1. The user requests a shell or the execution of a command, which begins the session mode.

In this mode, data is sent or received through the terminal on the client side, and the shell or command on the server side.
2. The user program terminates.

3. All X11 forwarding and TCP/IP forwarding is stopped. Any X11 connections and TCP/IP connections that already exist remain open.
4. The server sends the command `exit` to the client, and both sides exit.

---

## Configuring the Secure Shell

The characteristics of a Secure Shell session are controlled by configuration files, which can be overridden to a certain degree by options on the command line.

### Secure Shell Client Configuration

In most cases, the client-side characteristics of a Secure Shell session are governed by the system-wide configuration file, `/etc/ssh/ssh_config`, which is set up by the administrator. The settings in the system-wide configuration file can be overridden by the user's configuration in `$HOME/.ssh/config`. In addition, the user can override both configuration files on the command line.

The command line options are client requests and are permitted or denied on the server side by the `/etc/ssh/sshd_config` file (see `ssh_config(4)`). The configuration file keywords and command options are introduced in the following sections and are described in detail in the `ssh(1)`, `scp(1)`, `sftp(1)`, and `ssh_config(4)` man pages. Note that in the two user configuration files, the `Host` keyword indicates a host or wildcard expression to which all following keywords up to the next `Host` keyword apply.

### Host-Specific Parameters

If it is useful to have different Secure Shell characteristics for different local hosts, the administrator can define separate sets of parameters in the `/etc/ssh/ssh_config` file to be applied according to host or regular expression. This task is done by grouping entries in the file by `Host` keyword. If the `Host` keyword is not used, the entries in the client configuration file apply to whichever local host a user is working on.

### Client-Side Authentication Parameters

The authentication method is determined by setting one of the following keywords to "yes":

- `DSAAuthentication`

- PasswordAuthentication
- RhostsAuthentication
- RhostsRSAAuthentication

The keyword `UserRsh` specifies that the `rlogin` and `rsh` commands be used, probably due to no Secure Shell support.

The `Protocol` keyword sets the Secure Shell protocol version to `v1` or `v2`. You can specify both versions separated by a comma. The first version is tried and upon failure, the second version is used.

The `IdentityFile` keyword specifies an alternate file name to hold the user's private key.

The keyword `Cipher` specifies the `v1` encryption algorithm, which might be `blowfish` or `3des`. The keyword `Ciphers` specifies an order of preference for the `v2` encryption algorithms: `3des-cbc`, `blowfish-cbc`, and `aes128-cbc`. The commands `ssh` and `scp` have a `-c` option for specifying the encryption algorithm on the command line.

## Known Host File Parameters

The known host files (`/etc/ssh/ssh_known_hosts` and `$HOME/.ssh/known_hosts`) contain the public keys for all hosts with which the client can communicate by using Secure Shell. The `GlobalKnownHostsFile` keyword specifies an alternate file instead of `/etc/ssh/ssh_known_hosts`. The `UserKnownHostsFile` keyword specifies an alternate to `$HOME/.ssh/known_hosts`.

The `StrictHostKeyChecking` keyword requires new hosts to be added manually to the known hosts file, and refuses any host whose public key has changed or whose public key is not in the known hosts file. The keyword `CheckHostIP` enables the IP address for hosts in the known host files to be checked, in case a key has been changed due to DNS spoofing.

## Client-Side X11 Forwarding and Port Forwarding Parameters

The `LocalForward` keyword specifies a local TCP/IP port to be forwarded over a secure channel to a specified port on a remote host. The `GatewayPorts` keyword enables remote hosts to connect to local forwarded ports.

The command `ssh` enables port forwarding through these options:

- `-L`, which specifies the local port to be forwarded to the specified port on the remote host
- `-R`, which specifies a remote port to be forwarded to the local host and specified port



The `ForwardX11` keyword redirects X11 connections to the remote host with the `DISPLAY` environment variable set. The `XAuthLocation` keyword specifies the location of the `xauth(1)` program.

## Client-Side Connection and Other Parameters

The `NumberOfPasswordPrompts` keyword specifies how many times the user is prompted for a password before Secure Shell quits. The `ConnectionAttempts` keyword specifies how many tries (at one try per second) are made before Secure Shell either quits or falls back to `rsh` if the `FallBackToRsh` keyword is set.

The `Compression` keyword enables compression of transmitted data. The `CompressionLevel` keyword sets a level of 1 to 9, trading off between speed and amount of compression.

User specifies an alternate user name. `Hostname` specifies an alternate name for a remote host. `ProxyCommand` specifies an alternate command name for starting Secure Shell. Any command that can connect to your proxy server can be used. The command should read from its standard input and write to its standard output.

`Batchmode` disables password prompts, which is useful for scripts and other batch jobs.

`KeepAlive` enables messages to indicate network problems due to host crashes. `LogLevel` sets the verbosity level for `ssh` messages.

`EscapeChar` defines a single character that is used as a prefix for displaying special characters as plain text.

## Secure Shell Server Configuration

The server-side characteristics of a Secure Shell session are governed by the `/etc/ssh/sshd_config` file, which is set up by the administrator.

## Server-Side Authentication Parameters

Permitted authentication methods are indicated by these keywords:

- `DSAAuthentication`
- `PasswordAuthentication`
- `RhostsAuthentication`
- `RhostsRSAAuthentication`
- `RSAAuthentication`

`HostKey` and `HostDSAKey` identify files that hold host public keys when the default file name is not used. `KeyRegenerationInterval` defines how often the server key is regenerated.

`Protocol` specifies the version. `Ciphers` specifies the encryption algorithms for v2. `ServerKeyBits` defines the number of bits in the server's key.

## Ports and Forwarding Parameters

`AllowTCPForwarding` specifies whether TCP forwarding is permitted.

`GatewayPorts` allows remote hosts to connect to ports forwarded for the client. `Port` specifies the port number that `sshd` listens on. `ListenAddress` designates a specific local address that `sshd` listens to. If there is no `ListenAddress` specification, `sshd` listens to all addresses by default.

`X11Forwarding` allows X11 forwarding. `X11DisplayOffset` specifies the first display number that is available for forwarding. This keyword prevents `sshd` from interfering with real X11 servers. `XAuthLocation` specifies the location of the `xauth` program.

## Session Control Parameters

`KeepAlive` displays messages regarding broken connections and host crashes. `LogLevel` sets the verbosity level of messages from `sshd`. `SyslogFacility` provides a facility code for messages that are logged from `sshd`.

## Server Connection and Other Parameters

The `AllowGroups`, `AllowUsers`, `DenyGroups`, and `DenyUsers` keywords control which users can or cannot use `ssh`.

The `LoginGraceTime`, `MaxStartups`, `PermitRootLogin`, and `PermitEmptyPasswords` keywords set controls on users who are logging in. `StrictModes` causes `sshd` to check file modes and ownership of the user's files and home directory before login. `UseLogin` specifies whether `login` is used for interactive login sessions. Turning this keyword on should not be necessary and is not recommended for the Solaris environment.

`Subsystem` configures a file transfer daemon for using `sftp`.

---

## Maintaining Known Hosts on a Site-Wide Basis

Each host that needs to talk to another host securely must have the server's public key stored in the local host's `/etc/ssh/ssh_known_hosts` file. Although it is most convenient to update the `/etc/ssh/ssh_known_hosts` files by a script, this practice is heavily discouraged because it opens a major security vulnerability.

The `/etc/ssh/ssh_known_hosts` file should only be distributed by a secure mechanism as follows:

- Over a secure connection such as Secure Shell, IPsec, or kerberized `ftp` from a known and trusted machine
- At system install time

To avoid the possibility of an intruder gaining access by inserting bogus public keys into a `known_hosts` file, you should use the jumpstart server as the known and trusted source of the `ssh_known_hosts` file. The `ssh_known_hosts` file can be distributed during installation and by regularly running scripts on the individual hosts that pull in the latest version by using `scp`. This approach is secure because each host already has the public key of the jumpstart server.

---

## Secure Shell Files

The following table shows the important Secure Shell files and the suggested UNIX permissions.

**TABLE 5-1** Secure Shell Files

File Name	Description	Suggested Permissions and Owner
<code>/etc/ssh/sshd_config</code>	Contains configuration data for <code>sshd</code> , the Secure Shell daemon.	<code>-rw-r--r--</code> root
<code>/etc/ssh/ssh_host_key</code>	Contains the host private key.	<code>-rw-----</code> root
<code>/etc/ssh/ssh_host_key.pub</code>	Contains the host public key. Used to copy the host key to the local <code>known_hosts</code> file.	<code>-rw-r--r--</code> root

**TABLE 5-1** Secure Shell Files (Continued)

File Name	Description	Suggested Permissions and Owner
<code>/var/run/sshd.pid</code>	Contains the process ID of the Secure Shell daemon, <code>sshd</code> , which listens for connections (if there are multiple daemons, the file contains the last daemon that was started).	<code>rw-r--r-- root</code>
<code>\$(HOME)/.ssh/authorized_keys</code>	Lists the RSA keys that can be used with <code>v1</code> to log into the user's account, or the DSA and RSA keys that can be used with <code>v2</code> .	<code>-rw-rw-r-- johndoe</code>
<code>/etc/ssh/ssh_known_hosts</code>	Contains the host public keys for all hosts with which the client may communicate securely. The file should be prepared by the administrator.	<code>-rw-r--r-- root</code>
<code>\$(HOME)/.ssh/known_hosts</code>	Contains the host public keys for all hosts with which the client may communicate securely. The file is maintained automatically. Whenever the user connects with an unknown host, the remote host key is added to the file.	<code>-rw-r--r-- johndoe</code>
<code>/etc/nologin</code>	If this file exists, <code>sshd</code> refuses to let anyone except <code>root</code> log in. The contents are displayed to users who are attempting to log in.	<code>-rw-r--r-- root</code>
<code>\$(HOME)/.rhosts</code>	Contains the host-user name pairs that specifies the hosts to which the user can log in to without a password. The file is used Secure Shell, as well as by the <code>rlogind</code> and <code>rshd</code> daemons.	<code>-rw-r--r-- johndoe</code>
<code>\$(HOME)/.shosts</code>	Contains the host-user name pairs that specifies the hosts to which the user can log in to without a password using Secure Shell only.	<code>-rw-r--r-- johndoe</code>
<code>/etc/hosts.equiv</code>	Contains the hosts that are used in <code>.rhosts</code> authentication and Secure Shell authentication.	<code>-rw-r--r-- root</code>
<code>/etc/ssh/shosts.equiv</code>	Contains the hosts that are used in Secure Shell authentication.	<code>-rw-r--r-- root</code>
<code>\$(HOME)/.ssh/environment</code>	Used for initialization to make assignments at login.	<code>-rw----- johndoe</code>

**TABLE 5-1** Secure Shell Files (Continued)

File Name	Description	Suggested Permissions and Owner
<code>\$HOME/.ssh/rc</code>	Runs initialization routines before the user shell starts.	<code>-rw----- johndoe</code>
<code>/etc/ssh/sshrd</code>	Runs host-specific initialization routines that are specified by an administrator for all users.	<code>-rw-r--r-- root</code>

The following table summarizes the major Secure Shell commands.

**TABLE 5-2** Secure Shell Commands

Command	Description
<code>ssh</code>	A program for logging in to a remote machine and for executing commands on a remote machine. The command is intended to replace <code>rlogin</code> and <code>rsh</code> , and provide secure encrypted communications between two untrusted hosts over an insecure network. X11 connections and arbitrary TCP/IP ports can also be forwarded over the secure channel.
<code>sshd</code>	The daemon for Secure listens. This daemon listens for connections from clients and provides secure encrypted communications between two untrusted hosts over an insecure network.
<code>ssh-keygen</code>	Generates and manages authentication keys for <code>ssh</code> .
<code>ssh-agent</code>	A program that holds private keys that are used for public key authentication. <code>ssh-agent</code> is started at the beginning of an X-session or a login session, and all other windows or programs are started as clients to the <code>ssh-agent</code> program. Through the use of environment variables, the agent can be located and automatically used for authentication when users log in to other machines while using <code>ssh</code> .
<code>ssh-add</code>	Adds RSA or DSA identities (keys) to the authentication agent, <code>ssh-agent</code> .
<code>scp</code>	Securely copies files between hosts on a network by using <code>ssh</code> for data transfer. Unlike <code>rcp</code> , <code>scp</code> asks for passwords or passphrases (if they are needed for authentication).
<code>sftp</code>	An interactive file transfer program, similar to <code>ftp</code> , that performs all operations over an encrypted <code>ssh</code> transport. <code>sftp</code> connects and logs into the specified host name and then enters an interactive command mode.



## Introduction to SEAM

---

This chapter provides an introduction to Sun Enterprise Authentication Mechanism (SEAM).

- “What Is SEAM?” on page 71
- “How SEAM Works” on page 72
- “SEAM Security Services” on page 78
- “SEAM Releases” on page 78

---

### What Is SEAM?

SEAM is a client/server architecture that provides secure transactions over networks. SEAM offers strong user authentication, as well as data integrity and data privacy. *Authentication* guarantees that the identities of both the sender and the recipient of a network transaction are true. SEAM can also verify the validity of data being passed back and forth (*integrity*) and encrypt the data during transmission (*privacy*). Using SEAM, you can log on to other machines, execute commands, exchange data, and transfer files securely. Additionally, SEAM provides *authorization* services, which allows administrators to restrict access to services and machines. Moreover, as a SEAM user, you can regulate other people’s access to your account.

SEAM is a *single-sign-on* system, which means that you only need to authenticate yourself to SEAM once per session, and all subsequent transactions during the session are automatically secured. After SEAM has authenticated you, you do not need to authenticate yourself every time you use a SEAM-based command such as `kinit` or `klis`, or access data on an NFS file system. Thus, you do not have to send your password over the network, where it can be intercepted, each time you use these services.

SEAM is based on the Kerberos V5 network authentication protocol that was developed at the Massachusetts Institute of Technology (MIT). People who have used Kerberos V5 should therefore find SEAM very familiar. Since Kerberos V5 is a *de facto* industry standard for network security, SEAM promotes interoperability with other systems. In other words, because SEAM works with systems that use Kerberos V5, it allows for secure transactions even over heterogeneous networks. Moreover, SEAM provides authentication and security both between domains and within a single domain.

---

**Note** – Because SEAM is based on, and designed to interoperate with, Kerberos V5, this manual often uses the terms “Kerberos” and “SEAM” more or less interchangeably, for example, “Kerberos realm” or “SEAM-based utility.” Moreover, “Kerberos” and “Kerberos V5” are used interchangeably. The manual draws distinctions when necessary.

---

SEAM allows for flexibility in running Solaris applications. You can configure SEAM to allow both SEAM-based and non-SEAM-based requests for network services such as the NFS service. As a result, current Solaris applications still work even if they are running on systems on which SEAM is not installed. Of course, you can also configure SEAM to allow only SEAM-based network requests.

Additionally, applications do not have to remain committed to SEAM if other security mechanisms are developed. Because SEAM is designed to integrate modularly into the Generic Security Service (GSS) API, applications that make use of the GSS-API can utilize whichever security mechanism best suits its needs.

---

## How SEAM Works

The following is an overview of the SEAM authentication system. For a more detailed description, see “How the Authentication System Works” on page 189.

From the user’s standpoint, SEAM is mostly invisible after the SEAM session has been started. Initializing a SEAM session often involves no more than logging in and providing a Kerberos password.

The SEAM system revolves around the concept of a *ticket*. A ticket is a set of electronic information that serves as identification for a user or a service such as the NFS service. Just as your driver’s license identifies you and indicates what driving permissions you have, so a ticket identifies you and your network access privileges. When you perform a SEAM-based transaction (for example, if you `rlogin` in to another machine), you



transparently send a request for a ticket to a *Key Distribution Center*, or KDC. The KDC accesses a database to authenticate your identity and returns a ticket that grants you permission to access the other machine. “Transparently” means that you do not need to explicitly request a ticket.

Tickets have certain attributes associated with them. For example, a ticket can be *forwardable* (which means that it can be used on another machine without a new authentication process), or *postdated* (not valid until a specified time). How tickets are used (for example, which users are allowed to obtain which types of ticket) is set by *policies* that are determined when SEAM is installed or administered.

---

**Note** – You will frequently see the terms *credential* and *ticket*. In the greater Kerberos world, they are often used interchangeably. Technically, however, a credential is a ticket plus the *session key* for that session. This difference is explained in more detail in “Gaining Access to a Service Using SEAM” on page 190.

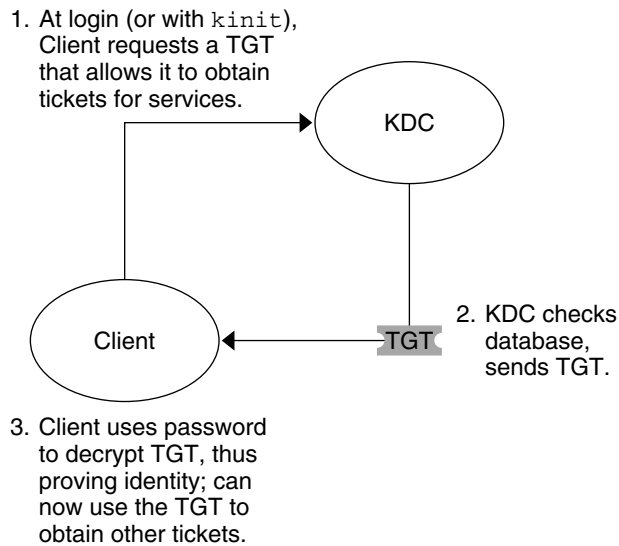
---

The following sections further explain the SEAM authentication process.

## Initial Authentication: the Ticket-Granting Ticket

Kerberos authentication has two phases: an initial authentication that allows for all subsequent authentications, and the subsequent authentications themselves.

The following figure shows how the initial authentication takes place.



TGT = Ticket-granting ticket  
 KDC = Key Distribution Center

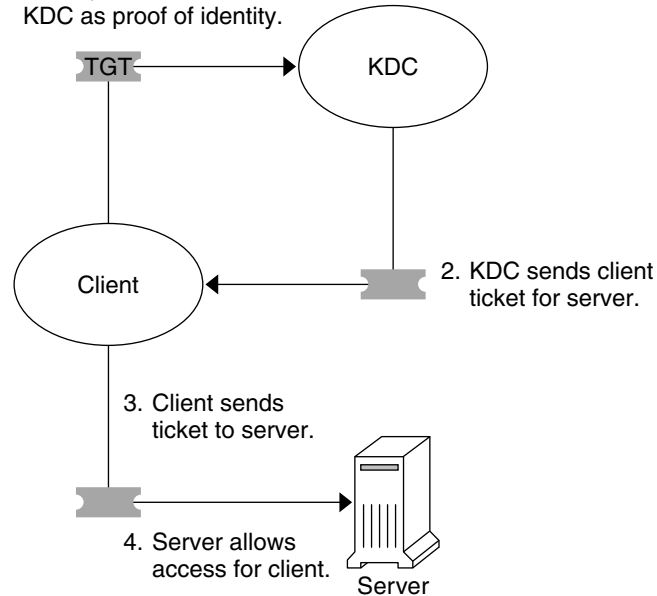
**FIGURE 6-1** Initial Authentication for SEAM Session

1. A client (a user, or a service such as NFS) begins a SEAM session by requesting a *ticket-granting ticket* (TGT) from the Key Distribution Center (KDC). This request is often done automatically at login.  
 A ticket-granting ticket is needed to obtain other tickets for specific services. Think of the ticket-granting ticket as similar to a passport. Like a passport, the ticket-granting ticket identifies you and allows you to obtain numerous “visas,” where the “visas” (tickets) are not for foreign countries but for remote machines or network services. Like passports and visas, the ticket-granting ticket and the other various tickets have limited lifetimes. The difference is that “Kerberized” commands notice that you have a passport and obtain the visas for you. You don’t have to perform the transactions yourself.
2. The KDC creates a ticket-granting ticket and sends it back, in encrypted form, to the client. The client decrypts the ticket-granting ticket by using the client’s password.
3. Now in possession of a valid ticket-granting ticket, the client can request tickets for all sorts of network operations for as long as the ticket-granting ticket lasts. This ticket usually lasts for a few hours. Each time the client performs a unique network operation, it requests a ticket for that operation from the KDC.

## Subsequent Authentications

After the client has received the initial authentication, each individual authentication follows the pattern that is shown in the following figure.

1. Client requests ticket for server; sends TGT to KDC as proof of identity.



TGT = Ticket-granting ticket  
KDC = Key Distribution Center

**FIGURE 6-2** Obtaining Access to a Service

1. The client requests a ticket for a particular service from the KDC by sending the KDC its ticket-granting ticket as proof of identity.
2. The KDC sends the ticket for the specific service to the client.

For example, suppose user *joe* wants to access an NFS file system that has been shared with krb5 authentication required. Since he is already authenticated (that is, he already has a ticket-granting ticket), as he attempts to access the files, the NFS client system automatically and transparently obtains a ticket from the KDC for the NFS service.

3. The client sends the ticket to the server.

When using the NFS service, the NFS client automatically and transparently sends the ticket for the NFS service to the NFS server.

4. The server allows the client access.

These steps make it appear that the server doesn't ever communicate with the KDC. The server does, though; it registers itself with the KDC, just as the first client does. For simplicity's sake, we have left that part out.

## Principals

A client in SEAM is identified by its *principal*. A principal is a unique identity to which the KDC can assign tickets. A principal can be a user, such as `joe`, or a service, such as `nfs`.

By convention, a principal name is divided into three parts: the *primary*, the *instance*, and the *realm*. A typical SEAM principal would be, for example, `joe/admin@ENG.EXAMPLE.COM`, where:

- `joe` is the primary. The primary can be a user name, as shown here, or a service, such as `nfs`. The primary can also be the word `host`, which signifies that this principal is a service principal that is set up to provide various network services.
- `admin` is the instance. An instance is optional in the case of user principals, but it is required for service principals. For example: if the user `joe` sometimes acts as a system administrator, he can use `joe/admin` to distinguish himself from his usual user identity. Likewise, if `joe` has accounts on two different hosts, he can use two principal names with different instances (for example, `joe/denver.example.com` and `joe/boston.example.com`). Notice that SEAM treats `joe` and `joe/admin` as two completely different principals.

In the case of a service principal, the instance is the fully qualified host name. `bigmachine.eng.example.com` is an example of such an instance so that the primary/instance might be, for example, `host/bigmachine.eng.example.com`.

- `ENG.EXAMPLE.COM` is the SEAM realm. Realms are discussed in "Realms" on page 76.

The following are all valid principal names:

- `joe`
- `joe/admin`
- `joe/admin@ENG.EXAMPLE.COM`
- `nfs/host.eng.example.com@ENG.EXAMPLE.COM`
- `host/eng.example.com@ENG.EXAMPLE.COM`

## Realms

A realm is a logical network, similar to a domain, which defines a group of systems under the same *master KDC*. Figure 6-3 shows how realms can relate to one another. Some realms are hierarchical (one realm being a superset of the other realm).

Otherwise, the realms are non-hierarchical (or "direct") and the mapping between the

two realms must be defined. A feature of SEAM is that it permits authentication across realms. Each realm only needs to have a principal entry for the other realm in its KDC. The feature is called cross-realm authentication.

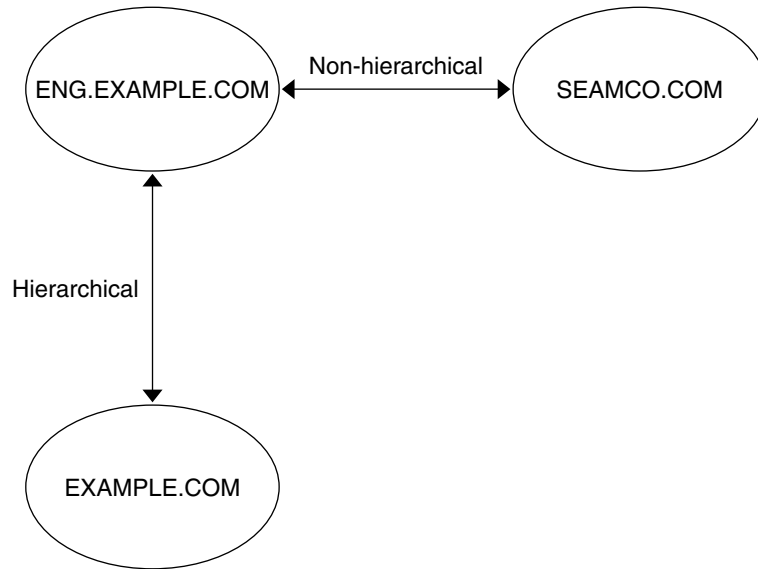



FIGURE 6-3 Realms

## Realms and Servers

Each realm must include a server that maintains the master copy of the principal database. This server is called the *master KDC server*. Additionally, each realm should contain at least one *slave KDC server*, which contains duplicate copies of the principal database. Both the master KDC server and the slave KDC server create tickets that are used to establish authentication.

Realms can also include *NFS servers*, which provide NFS services by using Kerberos authentication. If you have installed SEAM 1.0 or 1.0.1, the realm might include a SEAM network *application server*, which provides access to Kerberized applications (such as `ftp`, `telnet`, and `rsh`).

The following figure shows what a hypothetical realm might contain.



- Master KDC
- ▣ Clients
- Slave KDCs
- ▣ Application servers

FIGURE 6-4 A Typical Realm

---

## SEAM Security Services

In addition to providing secure authentication of users, SEAM provides two security services:

- **Integrity** – Just as authentication ensures that clients on a network are who they claim to be, integrity ensures that the data they send is valid and has not been tampered with during transit. Integrity is done through cryptographic checksumming of the data. Integrity also includes user authentication.
- **Privacy** – Privacy takes security a step further. Privacy not only includes verifying the integrity of transmitted data, but it encrypts the data before transmission, protecting it from eavesdroppers. Privacy authenticates users, as well.

---

## SEAM Releases

Components of the SEAM product have been included in four releases. The following table describes which components are included in each release. All components are described in the following sections.

**TABLE 6–1** SEAM Release Contents

Release Name	Contents
SEAM 1.0 in Solaris Easy Access Server (SEAS) 3.0	Full release of SEAM for the Solaris 2.6 and 7 releases
SEAM in the Solaris 8 release	SEAM client software only
SEAM 1.0.1 in the Solaris 8 Admin Pack	SEAM KDC and remote applications for the Solaris 8 release
SEAM in the Solaris 9 release	SEAM KDC and client software only

## SEAM 1.0 Components

Similar to the MIT distribution of Kerberos V5, SEAM includes the following:

- Key Distribution Center (KDC) (master):
  - Kerberos database administration daemon – `kadmind`
  - Kerberos ticket processing daemon – `krb5kdc`
- Slave KDCs
- Database administration programs – `kadmin` and `kadmin.local`
- Database propagation software – `kprop`
- User programs for obtaining, viewing, and destroying tickets – `kinit`, `klist`, `kdestroy` – and for changing your SEAM password – `kpasswd`
- Applications – `ftp`, `rcp`, `rlogin`, `rsh`, and `telnet` – and daemons for these applications – `ftpd`, `rlogind`, `rshd` and `telnetd`
- Administration utilities – `ktutil`, `kdb5_util`
- Several libraries

In addition, SEAM includes the following:

- SEAM Administration Tool (`gkadmin`) – Allows you to administer the KDC. This Java™ technology-based GUI allows an administrator to perform the tasks that are usually performed through the `kadmin` command.
- The Pluggable Authentication Module (PAM) – Allows applications to use various authentication mechanisms. PAM can be used to make login and logouts transparent to the user.
- A utility (`gsscred`) and a daemon (`gssd`) – These programs help map UNIX user IDs (UIDs) to principal names. These programs are needed because SEAM NFS servers use UNIX UIDs to identify users and not principal names, which are stored in a different format.
- The Generic Security Service Application Programming Interface (GSS-API) – Allows applications to use multiple security mechanisms without having to recompile the application every time a new mechanism is added. Because GSS-API

is machine-independent, it is appropriate for applications on the Internet. GSS-API provides applications with the ability to include the integrity and privacy security services, as well as authentication.

- The RPCSEC\_GSS Application Programming Interface (API) – Allows NFS services to use Kerberos authentication. RPCSEC\_GSS is a new security flavor that provides security services that are independent of the mechanisms being used. RPCSEC\_GSS sits “on top” of the GSS-API layer. Any pluggable GSS\_API-based security mechanism can be used by applications that use RPCSEC\_GSS.
- A preconfiguration procedure – Allows you to set the parameters for installing and configuring SEAM, which make SEAM installation automatic. This procedure is especially useful for multiple installations.
- Kernel modifications – Allows for faster performance.

## SEAM Components in the Solaris 8 Release

The Solaris 8 release included only the client-side portions of SEAM, so many components are not included. This product enables systems that run the Solaris 8 release to become SEAM clients without having to install SEAM separately. To use these capabilities, you must install a KDC that uses either SEAS 3.0 or the Solaris 8 Admin Pack, the MIT distribution, or Windows2000. The client-side components are not useful without a configured KDC to distribute tickets. The following components were included in this release:

- User programs for obtaining, viewing, and destroying tickets – `kinit`, `klist`, `kdestroy`.
- User program for changing your SEAM password – `kpasswd`.
- Key table administration utility – `ktutil`.
- Additions to the Pluggable Authentication Module (PAM) – Allows applications to use various authentication mechanisms. PAM can be used to make login and logouts transparent to the user.
- GSS\_API plug-ins – Provides Kerberos protocol and cryptographic support.
- NFS client and server support.

## SEAM 1.0.1 Components

The SEAM 1.0.1 release includes all components of the SEAM 1.0 release that are not already included in the Solaris 8 release. The components are as follows:

- Key Distribution Center (KDC) (master):
  - Kerberos database administration daemon – `kadmind`.
  - Kerberos ticket processing daemon – `krb5kdc`.
- Slave KDCs.



- Database administration programs – `kadmin` and `kadmin.local`.
- Database propagation software – `kprop`
- Applications – `ftp`, `rsh`, `rlogin`, `rsh`, and `telnet` – and daemons for these applications – `ftpd`, `rlogind`, `rshd` and `telnetd`.
- Administration utility – `kdb5_util`.
- SEAM Administration Tool (`gkadmin`) – Allows you to administer the KDC. This Java technology-based GUI allows an administrator to perform the tasks that are usually performed through the `kadmin` command.
- A preconfiguration procedure – Allows you to set the parameters for installing and configuring SEAM, which makes SEAM installation automatic. This procedure is especially useful for multiple installations.
- Several libraries.

## SEAM Components in the Solaris 9 Release

The Solaris 9 release includes all components of the SEAM 1.0 release, except for the applications and the preconfiguration procedure.



## Planning for SEAM

---

This chapter should be studied by administrators who are involved in the installation and maintenance of SEAM. The chapter discusses several installation and configuration issues that administrators must resolve before they install or configure SEAM.

This is a list of the issues that a system administrator or other knowledgeable support staff should resolve:

- “Realms” on page 84
- “Mapping Host Names Onto Realms” on page 85
- “Client and Service Principal Names” on page 85
- “Ports for the KDC and Admin Services” on page 86
- “Slave KDCs” on page 86
- “Database Propagation” on page 87
- “Clock Synchronization” on page 87

---

## Why Plan for SEAM?

Before you install SEAM, you must resolve several configuration issues. Although changing the configuration after the initial install is not impossible, it becomes more difficult with each new client that is added to the system. In addition, some changes require a full re-installation, so it is better to consider long-term goals when you planning your SEAM configuration.

---

## Realms

A realm is logical network, similar to a domain, which defines a group of systems that are under the same master KDC. As with establishing a DNS domain name, issues such as the realm name, the number and size of each realm, and the relationship of a realm to other realms for cross-realm authentication should be resolved before you configure SEAM.

### Realm Names

Realm names can consist of any ASCII string. Usually, the realm name is the same as your DNS domain name, in uppercase. This convention helps differentiate problems with SEAM from problems with the DNS namespace, while using a name that is familiar. If you do not use DNS or you choose to use a different string, then you can use any string. However, the configuration process requires more work. The use of realm names that follow the standard Internet naming structure is wise.

### Number of Realms

The number of realms that your installation requires depends on several factors:

- The number of clients to be supported. Too many clients in one realm makes administration more difficult and eventually requires that you split the realm. The primary factors that determine the number of clients that can be supported are as follows:
  - The amount of SEAM traffic that each client generates
  - The bandwidth of the physical network
  - The speed of the hosts

Since each installation will have different limitations, there is no rule for determining the maximum number of clients.

- How far apart the clients are. It might make sense to set up several small realms if the clients are in different geographic regions.
- The number of hosts that are available to be installed as KDCs. Each realm should have at least two KDC servers (master and slave).

### Realm Hierarchy

When you are configuring multiple realms for cross-realm authentication, you need to decide how to tie the realms together. You can establish a hierarchical relationship between the realms that provides automatic paths to the related domains. Of course,

all realms in the hierarchical chain must be configured properly. The automatic paths can ease the administration burden. However, if there are many levels of domains, you might not want to use the default path because it requires too many transactions.

You can also choose to establish the connection directly. A direct connection is most useful when too many levels exist between two hierarchical domains or when there is no hierarchal relationship. The connection must be defined in the `/etc/krb5/krb5.conf` file on all hosts that use the connection. So, some additional work is required. For an introduction, see “Realms” on page 76 and for the configuration procedures for multiple realms, see “Configuring Cross-Realm Authentication” on page 98.

---

## Mapping Host Names Onto Realms

The mapping of host names onto realm names is defined in the `domain_realm` section of the `krb5.conf` file. These mappings can be defined for a whole domain and for individual hosts, depending on the requirements. See the `krb5.conf(4)` man page for more information.

---

## Client and Service Principal Names

When you are using SEAM, it is strongly recommended that DNS services already be configured and running on all hosts. If DNS is used, it must be enabled on all systems or on none of them. If DNS is available, then the principal should contain the Fully Qualified Domain Name (FQDN) of each host. For example, if the host name is `boston`, the DNS domain name is `example.com`, and the realm name is `EXAMPLE.COM`, then the principal name for the host should be `host/boston.example.com@EXAMPLE.COM`. The examples in this book use the FQDN for each host.

For the principal names that include the FQDN of an host, it is important to match the string that describes the DNS domain name in the `/etc/resolv.conf` file. SEAM requires that the DNS domain name be in lowercase letters when you are entering the FQDN for a principal. The DNS domain name can include uppercase and lowercase letters, but only use lowercase letters when you are creating a host principal. For example, it doesn't matter if the DNS domain name is `example.com`, `Example.COM`, or any other variation. The principal name for the host would still be `host/boston.example.com@EXAMPLE.COM`.

SEAM can run without DNS services, but some key capabilities, such as the ability to communicate with other realms, will not work. If DNS is not configured, then a simple host name can be used as the instance name. In this case, the principal would be `host/boston@EXAMPLE.COM`. If DNS is enabled later, all host principals must be deleted and replaced in the KDC database.

---

## Ports for the KDC and Admin Services

By default, port 88 and port 750 are used for the KDC, and port 749 is used for the KDC administration daemon. Different port numbers can be used. However, if you change the port numbers, then the `/etc/services` and `/etc/krb5/krb5.conf` files must be changed on every client. In addition, the `/etc/krb5/kdc.conf` file on each KDC must be updated.

---

## Slave KDCs

Slave KDCs generate credentials for clients just as the master KDC does. The slave KDCs provide backup if the master becomes unavailable. Each realm should have at least one slave KDC. Additional slave KDCs might be required, depending on these factors:

- The number of physical segments in the realm. Normally, the network should be set up so that each segment can function, at least minimally, without the rest of the realm. To do so, a KDC must be accessible from each segment. The KDC in this instance could be either a master or a slave.
- The number of clients in the realm. By adding more slave KDC servers, you can reduce the load on the current servers.

It is possible to add too many slave KDCs. Remember that the KDC database must be propagated to each server, so the more KDC servers that are installed, the longer it can take to get the data updated throughout the realm. Also, since each slave retains a copy of the KDC database, more slaves increase the risk of a security breach.

In addition, one or more slave KDCs can easily be configured to be swapped with the master KDC. The advantage to following this procedure on at least one slave KDC is that if the master KDC fails for any reason, you will have a system preconfigured that will be easy to swap as the master KDC. For instructions on how to configure a swappable slave KDC, see “Swapping a Master KDC and a Slave KDC” on page 111.

---

## Database Propagation

The database that is stored on the master KDC must be regularly propagated to the slave KDCs. One of the first issues to resolve is how often to update the slave KDCs. The desire to have up-to-date information that is available to all clients needs to be weighed against the amount of time it takes to complete the update. For more information about database propagation, see “Administering the Kerberos Database” on page 114.

In large installations with many KDCs in one realm, it is possible for one or more slaves to propagate the data so that the process is done in parallel. This strategy reduces the amount of time that the update takes, but it also increases the level of complexity in administering the realm.

---

## Clock Synchronization

All hosts that participate in the Kerberos authentication system must have their internal clocks synchronized within a specified maximum amount of time. Known as *clock skew*, this feature provides another Kerberos security check. If the clock skew is exceeded between any of the participating hosts, requests are rejected.

One way to synchronize all the clocks is to use the Network Time Protocol (NTP) software. See “Synchronizing Clocks between KDCs and SEAM Clients” on page 109 for more information. Other ways of synchronizing the clocks are available, so the use of NTP is not required. However, some form of synchronization should be used to prevent access failures because of clock skew.

---

## Online Help URL

The online help URL is used by the SEAM Administration Tool, so the URL should be defined properly to enable the “Help Contents” menu to work. The HTML version of this manual can be installed on any appropriate server. Alternately, you can decide to use the collections at <http://docs.sun.com>.

The URL should point to the section titled “SEAM Administration Tool” in the “Administering Principals and Policies” chapter in this book. You can choose another HTML page, if another location is more appropriate.





## Configuring SEAM (Tasks)

---

This chapter provides configuration procedures for KDC servers, network application servers, NFS servers, and SEAM clients. Many of these procedures require superuser access, so they should be used by system administrators or advanced users. Cross-realm configuration procedures and other topics related to KDC servers are also covered.

- “Configuring KDC Servers” on page 90
- “Configuring Cross-Realm Authentication” on page 98
- “Configuring SEAM NFS Servers” on page 101
- “Configuring SEAM Clients” on page 106
- “Synchronizing Clocks between KDCs and SEAM Clients” on page 109
- “Swapping a Master KDC and a Slave KDC” on page 111
- “Administering the Kerberos Database” on page 114
- “Increasing Security” on page 120

---

## Configuring SEAM (Task Map)

Parts of the configuration process depend on other parts and must be done in a specific order. These procedures often establish services that are required to use SEAM. Other procedures are not dependent on any order, and can be done when appropriate. The following task map shows a suggested order for a SEAM installation.

**TABLE 8-1** First Steps: SEAM Configuration Order

Task	Description	For Instructions
1. Plan for your SEAM installation	Lets you resolve configuration issues before you start the software configuration process. Planning ahead saves you time and other resources in the long run.	Chapter 7
2. (Optional) Install NTP	Configures the Network Time Protocol (NTP) software, or another clock synchronization protocol. In order for SEAM to work properly, the clocks on all systems in the realm must be synchronized.	"Synchronizing Clocks between KDCs and SEAM Clients" on page 109
3. Configure the master KDC server	Configures and builds the master KDC server and database for a realm.	"How to Configure a Master KDC" on page 91
4. (Optional) Configure a slave KDC server	Configures and builds a slave KDC server for a realm.	"How to Configure a Slave KDC" on page 95
5. (Optional) Increase security on the KDC servers	Prevents security breaches on the KDC servers.	"How to Restrict Access to KDC Servers" on page 120
6. (Optional) Configure swappable KDC servers	Makes the task of swapping the master KDC and a slave KDC easier.	"How to Configure a Swappable Slave KDC" on page 111

Once the required steps have been completed, the following procedures can be used when required.

**TABLE 8-2** Next Steps: Additional SEAM Tasks

Task	Description	For Instructions
Configure cross-realm authentication	Enables communications from one realm to another realm.	"Configuring Cross-Realm Authentication" on page 98
Configure SEAM clients	Enables a client to use SEAM services.	"Configuring SEAM Clients" on page 106
Configure SEAM NFS server	Enables a server to share a file system that requires Kerberos authentication.	"Configuring SEAM NFS Servers" on page 101

## Configuring KDC Servers

After you install the SEAM software, you must configure the KDC servers. Configuring a master KDC and at least one slave KDC provides the service that issues credentials. These credentials are the basis for SEAM, so the KDCs must be installed before you attempt other tasks.

The most significant difference between a master KDC and a slave KDC is that only the master KDC can handle database administration requests. For instance, changing a password or adding a new principal must be done on the master KDC. These changes can then be propagated to the slave KDCs. Both the slave KDC and master KDC generate credentials. This feature provides redundancy in case the master KDC cannot respond.

## ▼ How to Configure a Master KDC

In this procedure, the following configuration parameters are used:

- Realm name = `EXAMPLE.COM`
- DNS domain name = `example.com`
- Master KDC = `kdc1.example.com`
- Slave KDC = `kdc2.example.com`
- admin principal = `kws/admin`
- Online help URL =  
`http://denver:8888/ab2/coll.384.1/SEAM/@AB2PageView/6956`

---

**Note** – Adjust the URL to point to the “SEAM Administration Tool” section, as described in “Online Help URL” on page 87.

---

### 1. Complete the prerequisites for configuring a master KDC.

This procedure requires that DNS must be running. For specific naming instructions if this master is to be swappable, see “Swapping a Master KDC and a Slave KDC” on page 111.

### 2. Become superuser on the master KDC.

### 3. Edit the Kerberos configuration file (`krb5.conf`).

You need to change the realm names and the names of the servers. See the `krb5.conf(4)` man page for a full description of this file.

```
kdc1 # cat /etc/krb5/krb5.conf
[libdefaults]
    default_realm = EXAMPLE.COM

[realms]
    EXAMPLE.COM = {
        kdc = kdc1.example.com
        kdc = kdc2.example.com
        admin_server = kdc1.example.com
    }
```

```

[domain_realm]
    .example.com = EXAMPLE.COM
#
# if the domain name and realm name are equivalent,
# this entry is not needed
#
[logging]
    default = FILE:/var/krb5/kdc.log
    kdc = FILE:/var/krb5/kdc.log

[appdefaults]
    gkadmin = {
        help_url = http://denver:8888/ab2/coll.384.1/SEAM/@AB2PageView/6956
    }

```

In this example, the lines for `domain_realm`, `kdc`, `admin_server`, and all `domain_realm` entries were changed. In addition, the line that defines the `help_url` was edited.

#### 4. Edit the KDC configuration file (`kdc.conf`).

You need to change the realm name. See the `kdc.conf(4)` man page for a full description of this file.

```

kdc1 # cat /etc/krb5/kdc.conf
[kdcdefaults]
    kdc_ports = 88,750

[realms]
    EXAMPLE.COM= {
        profile = /etc/krb5/krb5.conf
        database_name = /var/krb5/principal
        admin_keytab = /etc/krb5/kadm5.keytab
        acl_file = /etc/krb5/kadm5.acl
        kadmind_port = 749
        max_life = 8h 0m 0s
        max_renewable_life = 7d 0h 0m 0s
    }

```

In this example, the realm name definition in the `realms` section was changed.

#### 5. Create the KDC database by using the `kdb5_util` command.

The `kdb5_util` command creates the KDC database. Also, when used with the `-s` option, this command creates a stash file that is used to authenticate the KDC to itself before the `kadmind` and `krb5kdc` daemons are started.

```

kdc1 # /usr/sbin/kdb5_util create -r EXAMPLE.COM -s
Initializing database '/var/krb5/principal' for realm 'EXAMPLE.COM'
master key name 'K/M@EXAMPLE.COM'
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.
Enter KDC database master key:      <type the key>
Re-enter KDC database master key to verify:  <type it again>

```

The `-r` option followed by the realm name is not required if the realm name is equivalent to the domain name in the server's name space.

## 6. Edit the Kerberos access control list file (`kadm5.ac1`).

Once populated, the `/etc/krb5/kadm5.ac1` file should contain all principal names that are allowed to administer the KDC. The first entry that is added might look similar to the following:

```
kws/admin@EXAMPLE.COM *
```

This entry gives the `kws/admin` principal in the `EXAMPLE.COM` realm the ability to modify principals or policies in the KDC. The default installation includes an asterisk (\*) to match all admin principals. This default could be a security risk, so it is more secure to include a list of all of the admin principals. See the `kadm5.ac1(4)` man page for more information.

## 7. Start the `kadmin.local` command.

The next sub-steps create principals that are used by SEAM.

```
kdc1 # /usr/sbin/kadmin.local
kadmin.local:
```

### a. Add administration principals to the database.

You can add as many admin principals as you need. You must add at least one admin principal to complete the KDC configuration process. For this example, a `kws/admin` principal is added. You can substitute an appropriate principal name instead of "kws."

```
kadmin.local: addprinc kws/admin
Enter password for principal kws/admin@EXAMPLE.COM: <type the password>
Re-enter password for principal kws/admin@EXAMPLE.COM: <type it again>
Principal "kws/admin@EXAMPLE.COM" created.
kadmin.local:
```

### b. Create a keytab file for the `kadmin` service.

This command sequence creates a special keytab file with principal entries for `kadmin` and `changepw`. These principals are needed for the `kadmin` service. Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab kadmin/kdc1.example.com
Entry for principal kadmin/kdc1.example.com with kvno 3, encryption type DES-CBC-CRC
added to keytab WRFILE:/etc/krb5/kadm5.keytab.
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab changepw/kdc1.example.com
Entry for principal changepw/kdc1.example.com with kvno 3, encryption type DES-CBC-CRC
added to keytab WRFILE:/etc/krb5/kadm5.keytab.
kadmin.local:
```

### c. Quit `kadmin.local`.

You have added all of the required principals for the next steps.

```
kadmin.local: quit
```

## 8. Start the Kerberos daemons.

```
kdc1 # /etc/init.d/kdc start
kdc1 # /etc/init.d/kdc.master start
```

## 9. Start kadmin.

At this point, you can add principals by using the SEAM Administration Tool. To do so, you must log on with one of the admin principal names that you created earlier in this procedure. However, the following command-line example is shown for simplicity.

```
kdc1 # /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
kadmin:
```

### a. Create the master KDC host principal.

The host principal is used by Kerberized applications (such as `klist` and `kprop`). Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kadmin: addprinc -randkey host/kdc1.example.com
Principal "host/kdc1.example.com@EXAMPLE.COM" created.
kadmin:
```

### b. (Optional) Create the master KDC root principal.

This principal is used for authenticated NFS-mounting. So, the principal might not be necessary on a master KDC. Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kadmin: addprinc root/kdc1.example.com
Enter password for principal root/kdc1.example.com@EXAMPLE.COM: <type the password>
Re-enter password for principal root/kdc1.example.com@EXAMPLE.COM: <type it again>
Principal "root/kdc1.example.com@EXAMPLE.COM" created.
kadmin:
```

### c. Add the master KDC's host principal to the master KDC's keytab file.

Adding the host principal to the keytab file allows this principal to be used automatically.

```
kadmin: ktadd host/kdc1.example.com
kadmin: Entry for principal host/kdc1.example.com with
kvno 3, encryption type DES-CBC-CRC added to keytab
WRFILE:/etc/krb5/krb5.keytab
kadmin:
```

### d. Quit kadmin.

```
kadmin: quit
```

## 10. Add an entry for each KDC into the propagation configuration file (`kpropd.acl`).

See the `kprop(1M)` man page for a full description of this file.

```
kdc1 # cat /etc/krb5/kpropd.acl
host/kdc1.example.com@EXAMPLE.COM
host/kdc2.example.com@EXAMPLE.COM
```

### 11. (Optional) Synchronize the master KDCs clock by using NTP or another clock synchronization mechanism.

It is not required to install and use the Network Time Protocol (NTP). However, every clock must be within the default time that is defined in the `libdefaults` section of the `krb5.conf` file in order for authentication to succeed. See “Synchronizing Clocks between KDCs and SEAM Clients” on page 109 for information about NTP.

## ▼ How to Configure a Slave KDC

In this procedure, a new slave KDC named `kdc3` is configured. This procedure uses the following configuration parameters:

- Realm name = `EXAMPLE.COM`
- DNS domain name = `example.com`
- Master KDC = `kdc1.example.com`
- Slave KDC = `kdc2.example.com` and `kdc3.example.com`
- admin principal = `kws/admin`
- Online help URL =  
`http://denver:8888/ab2/coll.384.1/SEAM/@AB2PageView/6956`

---

**Note** – Adjust the URL to point to the “SEAM Administration Tool” section, as described in the “Online Help URL” on page 87.

---

### 1. Complete the prerequisites for configuring a slave KDC.

The master KDC must be configured. For specific instructions if this slave is to be swappable, see “Swapping a Master KDC and a Slave KDC” on page 111.

### 2. On the master KDC, become superuser.

### 3. On the master KDC, start `kadmin`.

You must log on with one of the admin principal names that you created when you configure the master KDC.

```
kdc1 # /usr/sbin/kadmin -p kws/admin
Enter password: <Enter kws/admin password>
kadmin:
```

### a. On the master KDC, add slave host principals to the database, if not already done.

In order for the slave to function, it must have a host principal. Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kadmin: addprinc -randkey host/kdc3.example.com
Principal "host/kdc3@EXAMPLE.COM" created.
kadmin:
```

**b. (Optional) On the master KDC, create the slave KDC root principal.**

This principal is only needed if the slave will be NFS-mounting an authenticated file system. Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kadmin: addprinc root/kdc3.example.com
Enter password for principal root/kdc3.example.com@EXAMPLE.COM: <type the password>
Re-enter password for principal root/kdc3.example.com@EXAMPLE.COM: <type it again>
Principal "root/kdc3.example.com@EXAMPLE.COM" created.
kadmin:
```

**c. Quit kadmin.**

```
kadmin: quit
```

**4. On the master KDC, edit the Kerberos configuration file (`krb5.conf`).**

You need to add an entry for each slave. See the `krb5.conf(4)` man page for a full description of this file.

```
kdc1 # cat /etc/krb5/krb5.conf
[libdefaults]
    default_realm = EXAMPLE.COM

[realms]
    EXAMPLE.COM = {
        kdc = kdc1.example.com
        kdc = kdc2.example.com
        kdc = kdc3.example.com
        admin_server = kdc1.example.com
    }

[domain_realm]
    .example.com = EXAMPLE.COM
#
# if the domain name and realm name are equivalent,
# this entry is not needed
#
[logging]
    default = FILE:/var/krb5/kdc.log
    kdc = FILE:/var/krb5/kdc.log

[appdefaults]
    gkadmin = {
        help_url = http://denver:8888/ab2/coll.384.1/SEAM/@AB2PageView/6956
```



5. On the master KDC, add an entry for each slave KDC into the database propagation configuration file (`kpropd.acl`).

See the `kprop(1M)` man page for a full description of this file.

```
kdc1 # cat /etc/krb5/kpropd.acl
host/kdc1.example.com@EXAMPLE.COM
host/kdc2.example.com@EXAMPLE.COM
host/kdc3.example.com@EXAMPLE.COM
```

6. On all slave KDCs, copy the KDC administration files from the master KDC server.

This step needs to be followed on all slave KDCs, since the master KDC server has updated information that each KDC server needs. You can use `ftp` or a similar transfer mechanism to grab copies of the following files from the master KDC:

- `/etc/krb5/krb5.conf`
- `/etc/krb5/kdc.conf`
- `/etc/krb5/kpropd.acl`

7. On the new slave, add the slave's host principal to the slave's keytab file by using `kadmin`.

You must log on with one of the admin principal names that you created when you configure the master KDC. This entry allows `kprop` and other Kerberized applications to function. Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kdc3 # /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
kadmin: ktadd host/kdc3.example.com
kadmin: Entry for principal host/kdc3.example.com with
kvno 3, encryption type DES-CBC-CRC added to keytab
WRFILE:/etc/krb5/krb5.keytab
kadmin: quit
```

8. On the master KDC, add slave KDC names to the cron job, which automatically runs the backups, by running `crontab -e`.

Add the name of each slave KDC server at the end of the `kprop_script` line.

```
10 3 * * * /usr/lib/krb5/kprop_script kdc2.example.com kdc3.example.com
```

You might also want to change the time of the backups. This configuration starts the backup process every day at 3:10 AM.

9. On the master KDC, back up and propagate the database by using `kprop_script`.

If a backup copy of the database is already available, it is not necessary to complete another backup. See "How to Manually Propagate the Kerberos Database to the Slave KDCs" on page 118 for further instructions.

```
kdc1 # /usr/lib/krb5/kprop_script kdc3.example.com
Database propagation to kdc3.example.com: SUCCEEDED
```

10. On the new slave, create a stash file by using `kdb5_util`.

```
kdc3 # /usr/sbin/kdb5_util stash
kdb5_util: Cannot find/read stored master key while reading master key
kdb5_util: Warning: proceeding without master key
```

```
Enter KDC database master key: <type the key>
```

**11. (Optional) On the new slave KDC, synchronize the master KDCs clock by using NTP or another clock synchronization mechanism.**

It is not required to install and use the Network Time Protocol (NTP). However, every clock must be within the default time that is defined in the `libdefaults` section of the `krb5.conf` file in order for authentication to succeed. See “Synchronizing Clocks between KDCs and SEAM Clients” on page 109 for information about NTP.

**12. On the new slave, start the KDC daemon (`krb5kdc`).**

```
kdc3 # /etc/init.d/kdc start
```

---

## Configuring Cross-Realm Authentication

You have several ways of linking realms together so that users in one realm can be authenticated in another realm. Normally, this cross-realm authentication is accomplished by establishing a secret key that is shared between the two realms. The relationship of the realms can be either hierarchal or directional (see “Realm Hierarchy” on page 84).

### ▼ How to Establish Hierarchical Cross-Realm Authentication

The example in this procedure uses two realms, `ENG.EAST.EXAMPLE.COM` and `EAST.EXAMPLE.COM`. Cross-realm authentication will be established in both directions. This procedure must be completed on the master KDC in both realms.

**1. Complete the prerequisites for establishing hierarchical cross-realm authentication.**

The master KDC for each realm must be configured. To fully test the authentication process, several clients or slave KDCs must be installed.

**2. Become superuser on the first master KDC.**

**3. Create ticket-granting ticket service principals for the two realms.**

You must log on with one of the `admin` principal names that was created when you configured the master KDC.

```
# /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
```

```

kadmin: addprinc krbtgt/ENG.EAST.EXAMPLE.COM@EAST.EXAMPLE.COM
Enter password for principal krbtgt/ENG.EAST.EXAMPLE.COM@EAST.EXAMPLE.COM: <type the
                                                                    password>

kadmin: addprinc krbtgt/EAST.EXAMPLE.COM@ENG.EAST.EXAMPLE.COM
Enter password for principal krbtgt/EAST.EXAMPLE.COM@ENG.EAST.EXAMPLE.COM: <type the
                                                                    password>

kadmin: quit

```

---

**Note** – The password that is entered for each service principal must be identical in both KDCs. Thus, the password for the service principal `krbtgt/ENG.EAST.EXAMPLE.COM@EAST.EXAMPLE.COM` must be the same in both realms.

---

**4. Add entries to the Kerberos configuration file to define domain names for every realm (`krb5.conf`).**

```

# cat /etc/krb5/krb5.conf
[libdefaults]
.
.
[domain_realm]
    .eng.east.example.com = ENG.EAST.EXAMPLE.COM
    .east.example.com = EAST.EXAMPLE.COM

```

In this example, domain names for the `ENG.EAST.EXAMPLE.COM` and `EAST.EXAMPLE.COM` realms are defined. It is important to include the subdomain first, since the file is searched top down.

**5. Copy the Kerberos configuration file to all clients in this realm.**

In order for cross-realm authentication to work, all systems (including slave KDCs and other servers) must have the new version of the Kerberos configuration file (`/etc/krb5/krb5.conf`) installed.

**6. Repeat these steps in the second realm.**

## ▼ How to Establish Direct Cross-Realm Authentication

The example in this procedure uses two realms: `ENG.EAST.EXAMPLE.COM` and `SALES.WEST.EXAMPLE.COM`. Cross-realm authentication will be established in both directions. This procedure must be completed on the master KDC in both realms.

**1. Complete the prerequisites for establishing direct cross-realm authentication.**

The master KDC for each realm must be configured. To fully test the authentication process, several clients or slave KDCs must be installed.

**2. Become superuser on one of the master KDC servers.**

### 3. Create ticket-granting ticket service principals for the two realms.

You must log on with one of the admin principal names that was created when you configured the master KDC.

```
# /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
kadmin: addprinc krbtgt/ENG.EAST.EXAMPLE.COM@SALES.WEST.EXAMPLE.COM
Enter password for principal
krtgt/ENG.EAST.EXAMPLE.COM@SALES.WEST.EXAMPLE.COM: <type the password>
kadmin: addprinc krbtgt/SALES.WEST.EXAMPLE.COM@ENG.EAST.EXAMPLE.COM
Enter password for principal
krtgt/SALES.WEST.EXAMPLE.COM@ENG.EAST.EXAMPLE.COM: <type the password>
kadmin: quit
```

---

**Note** – The password that is entered for each service principal must be identical in both KDCs. Thus, the password for the service principal `krbtgt/ENG.EAST.EXAMPLE.COM@SALES.WEST.EXAMPLE.COM` must be the same in both realms.

---

### 4. Add entries in the Kerberos configuration file to define the direct path to the remote realm (`krb5.conf`).

This example shows the clients in the `ENG.EAST.EXAMPLE.COM` realm. You would need to swap the realm names to get the appropriate definitions in the `SALES.WEST.EXAMPLE.COM` realm.

```
# cat /etc/krb5/krb5.conf
[libdefaults]
.
.
.
[capaths]
    ENG.EAST.EXAMPLE.COM = {
        SALES.WEST.EXAMPLE.COM = .
    }

    SALES.WEST.EXAMPLE.COM = {
        ENG.EAST.EXAMPLE.COM = .
    }
```

### 5. Copy the Kerberos configuration file to all clients in the current realm.

In order for cross-realm authentication to work, all systems (including slave KDCs and other servers) must have the new version of the Kerberos configuration file (`krb5.conf`) installed.

### 6. Repeat these steps for the second realm.

---

# Configuring SEAM NFS Servers

NFS services use UNIX user IDs (UIDs) to identify a user and cannot directly use principals. To translate the principal to a UID, a credential table that maps user principals to UNIX UIDs must be created. The procedures in this section focus on the tasks that are necessary to configure a SEAM NFS server, to administer the credential table, and to initiate Kerberos security modes for NFS-mounted file systems. The following task map describes the tasks that are covered in this section.

**TABLE 8-3** Configuring SEAM NFS Servers (Task Map)

Task	Description	For Instructions
Configure a SEAM NFS server	Enables a server to share a file system that requires Kerberos authentication.	"How to Configure SEAM NFS Servers" on page 101
Create a credential table	Generates a credential table.	"How to Create a Credential Table" on page 103
Change the credential table that maps user principals to UNIX UIDs	Updates information in the credential table.	"How to Add a Single Entry to the Credential Table" on page 103
Share a file system with Kerberos authentication	Shares a file system with security modes so that Kerberos authentication is required.	"How to Set Up a Secure NFS Environment With Multiple Kerberos Security Modes" on page 104

## ▼ How to Configure SEAM NFS Servers

In this procedure, the following configuration parameters are used:

- Realm name = `EXAMPLE.COM`
- DNS domain name = `example.com`
- NFS server = `denver.example.com`
- admin principal = `kws/admin`

### 1. Complete the prerequisites for configuring a SEAM NFS server.

The master KDC must be configured. To fully test the process, you need several clients.

### 2. (Optional) Install the NTP client or other clock synchronization mechanism.

It is not required to install and use the Network Time Protocol (NTP). However, every clock must be within the default time that is defined in the `libdefaults` section of the `krb5.conf` file in order for authentication to succeed. See "Synchronizing Clocks between KDCs and SEAM Clients" on page 109 for information about NTP.

### 3. Start kadmin.

You can use the SEAM Administration Tool to add a principal, as explained in “How to Create a New Principal” on page 146. To do so, you must log on with one of the admin principal names that you created when you configured the master KDC. However, the following example shows how to add the required principals by using the command line.

```
denver # /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
kadmin:
```

#### a. Create the server’s NFS service principal.

Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```
kadmin: addprinc -randkey nfs/denver.example.com
Principal "nfs/denver.example.com" created.
kadmin:
```

#### b. (Optional) Create a root principal for the NFS server.

```
kadmin: addprinc root/denver.example.com
Enter password for principal root/denver.example.com@EXAMPLE.COM: <type the password>
Re-enter password for principal root/denver.example.com@EXAMPLE.COM: <type it again>
Principal "root/denver.example.com@EXAMPLE.COM" created.
kadmin:
```

#### c. Add the server’s NFS service principal to the server’s keytab file.

```
kadmin: ktadd nfs/denver.example.com
kadmin: Entry for principal nfs/denver.example.com with
kvno 3, encryption type DES-CBC-CRC added to keytab
WRFILE: /etc/krb5/krb5.keytab
kadmin:
```

#### d. Quit kadmin.

```
kadmin: quit
```

### 4. Create the gsscred table.

See “How to Create a Credential Table” on page 103 for more information.

### 5. Share the NFS file system with Kerberos security modes.

See “How to Set Up a Secure NFS Environment With Multiple Kerberos Security Modes” on page 104 for more information.

### 6. On each client, authenticate both the user principal and the root principal.

See “Setting Up Root Authentication to Mount NFS File Systems” on page 109 for more information.

## ▼ How to Create a Credential Table

The `gsscred` credential table is used by an NFS server to map SEAM principals to a UID. In order for NFS clients to mount file systems from an NFS server with Kerberos authentication, this table must be created or made available.

1. **Edit `/etc/gss/gsscred.conf` and change the mechanism.**

Change the mechanism to `files`.

2. **Create the credential table by using `gsscred`.**

```
# gsscred -m kerberos_v5 -a
```

The `gsscred` command gathers information from all sources that are listed with the `passwd` entry in the `/etc/nsswitch.conf` file. You might need to temporarily remove the `files` entry, if you do not want the local password entries included in the credential table. See the `gsscred(1M)` man page for more information.

## ▼ How to Add a Single Entry to the Credential Table

This procedure requires that the `gsscred` table has already been created on the NFS server.

1. **Become superuser on a NFS server.**

2. **Add an entry to the table by using `gsscred`.**

```
# gsscred -m mech [ -n name [ -u uid ] ] -a
```

<i>mech</i>	Defines the security mechanism to be used.
<i>name</i>	Defines the principal name for the user, as defined in the KDC.
<i>uid</i>	Defines the UID for the user, as defined in the password database.
<i>-a</i>	Adds the UID to principal name mapping.

### Example—Adding a Single Entry to the Credential Table

In the following example, an entry is added for the user named `sandy`, which is mapped to UID 3736. The UID is pulled from the password file if it is not included on the command line.

```
# gsscred -m kerberos_v5 -n sandy -u 3736 -a
```

## ▼ How to Set Up a Secure NFS Environment With Multiple Kerberos Security Modes

### 1. Become superuser on the NFS server.

### 2. Verify that there is a NFS service principal in the keytab file.

The `klist` command reports if there is a keytab file and displays the principals. If the results show that there is no keytab file or that there is no NFS service principal, you need to verify the completion of all of the steps in “How to Configure SEAM NFS Servers” on page 101.

```
# klist -k
Keytab name: FILE:/etc/krb5/krb5.keytab
KVNO Principal
-----
    3 nfs/denver.example.com@EXAMPLE.COM
```

### 3. Enable Kerberos security modes in the `/etc/nfssec.conf` file.

Edit the `/etc/nfssec.conf` file and remove the “#” from in front of the Kerberos security modes.

```
# cat /etc/nfssec.conf
.
.
#
# Uncomment the following lines to use Kerberos V5 with NFS
#
krb5          390003  kerberos_v5   default -      # RPCSEC_GSS
krb5i        390004  kerberos_v5   default integrity # RPCSEC_GSS
krb5p        390005  kerberos_v5   default privacy  # RPCSEC_GSS
```

### 4. Edit the `/etc/dfs/dfstab` file and add the `sec=` option with the required security modes to the appropriate entries.

```
share -F nfs -o sec=mode file-system
```

<i>mode</i>	Specifies the security modes to be used when sharing. When using multiple security modes, the first mode in the list is used as the default by the automounter.
<i>file-system</i>	Defines the path to the file system to be shared.

All clients that attempt to access files from the named file system require Kerberos authentication. To access files, both the user principal and the `root` principal on the NFS client should be authenticated.

### 5. Make sure that the NFS service is running on the server.



If this command is the first `share` command or set of `share` commands that you have initiated, it is likely that the NFS daemons are not running. The following commands kill the daemons and restart them.

```
# /etc/init.d/nfs.server stop
# /etc/init.d/nfs.server start
```

**6. (Optional) If the automounter is being used, edit the `auto_master` database to select a security mode other than the default.**

You need not follow this procedure if you are not using the automounter to access the file system or if the default selection for the security mode is acceptable.

```
file-system auto_home -nosuid,sec=mode
```

**7. (Optional) Manually issue the `mount` command to access the file system by using a non-default mode.**

Alternatively, you could use the `mount` command to specify the security mode, but this alternative does not take advantage of the automounter:

```
# mount -F nfs -o sec=mode file-system
```

## Example—Sharing a File System With One Kerberos Security Mode

In this example, the `dfstab` file line means that Kerberos authentication must succeed before any files can be accessed through the NFS service.

```
# grep krb /etc/dfs/dfstab
share -F nfs -o sec=krb5 /export/home
```

## Example—Sharing a File System With Multiple Kerberos Security Modes

In this example, all three Kerberos security modes have been selected. If no security mode is specified when a `mount` request is made, the first mode that is listed is used on all NFS V3 clients (in this case, `krb5`). See the `nfssec.conf(4)` man page for more information.

```
# grep krb /etc/dfs/dfstab
share -F nfs -o sec=krb5:krb5i:krb5p /export/home
```

---

# Configuring SEAM Clients

SEAM clients include any host, not a KDC server, on the network that needs to use SEAM services. This section provides a procedure for installing a SEAM client, as well as specific information about using root authentication to mount NFS file systems.

## ▼ How to Configure a SEAM Client

In this procedure, the following configuration parameters are used:

- Realm name = `EXAMPLE.COM`
- DNS domain name = `example.com`
- Master KDC = `kdc1.example.com`
- Slave KDC = `kdc2.example.com`
- Client = `client.example.com`
- admin principal = `kws/admin`
- User principal = `mre`
- Online help URL =  
`http://denver:8888/ab2/coll.384.1/SEAM/@AB2PageView/6956`

---

**Note** – Adjust the URL to point to the “SEAM Administration Tool” section, as described in the “Online Help URL” on page 87.

---

### 1. Become superuser.

### 2. Edit the Kerberos configuration file (`krb5.conf`).

To change the file from the SEAM default version, you need to change the realm names and the names of the servers. You also need to identify the path to the help files for `gkadmin`.

```
kdc1 # cat /etc/krb5/krb5.conf
[libdefaults]
    default_realm = EXAMPLE.COM

[realms]
    EXAMPLE.COM = {
        kdc = kdc1.example.com
        kdc = kdc2.example.com
        admin_server = kdc1.example.com
    }
```

```

[domain_realm]
    .example.com = EXAMPLE.COM
#
# if the domain name and realm name are equivalent,
# this entry is not needed
#
[logging]
    default = FILE:/var/krb5/kdc.log
    kdc = FILE:/var/krb5/kdc.log

[appdefaults]
    gkadmin = {
        help_url = http://denver:8888/ab2/coll.384.1/SEAM/@AB2PageView/6956

```

**3. (Optional) Synchronize the client’s clock with the master KDC’s clock by using NTP or another clock synchronization mechanism.**

It is not required to install and use the Network Time Protocol (NTP). However, every clock must be within the default time that is defined in the `libdefaults` section of the `krb5.conf` file in order for authentication to succeed. See “Synchronizing Clocks between KDCs and SEAM Clients” on page 109 for information about NTP.

**4. (Optional) Create a user principal if a user principal does not already exist.**

You need to create a user principal only if the user associated with this host does not have a principal assigned already. See “How to Create a New Principal” on page 146 for instructions on using the SEAM Administration Tool. The following is a command-line example.

```

client1 # /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
kadmin: addprinc mre
Enter password for principal mre@EXAMPLE.COM: <type the password>
Re-enter password for principal mre@EXAMPLE.COM: <type it again>
kadmin:

```

**5. Create a root principal.**

Note that when the principal instance is a host name, the FQDN must be entered in lowercase letters, regardless of the case of the domainname in the `/etc/resolv.conf` file.

```

kadmin: addprinc root/client1.example.com
Enter password for principal root/client1.example.com@EXAMPLE.COM: <type the password>
Re-enter password for principal root/client1.example.com@EXAMPLE.COM: <type it again>
kadmin: quit

```

**6. (Optional) To use Kerberos with NFS, enable Kerberos security modes in the `/etc/nfssec.conf` file.**

Edit the `/etc/nfssec.conf` file and remove the “#” from in front of the Kerberos security modes.

```

# cat /etc/nfssec.conf
.
.

```

```
#
# Uncomment the following lines to use Kerberos V5 with NFS
#
krb5          390003  kerberos_v5      default -          # RPCSEC_GSS
krb5i        390004  kerberos_v5      default integrity  # RPCSEC_GSS
krb5p        390005  kerberos_v5      default privacy   # RPCSEC_GSS
```

**7. (Optional) If you want a user on the SEAM client to automatically mount Kerberized NFS file systems that use Kerberos authentication, you must authenticate the root user.**

This process is done most securely by using the `kinit` command. However, users will need to use `kinit` as `root` every time they need to mount a file system that is secured by Kerberos. You can choose to use a keytab file instead. For detailed information about the keytab file requirement, see “Setting Up Root Authentication to Mount NFS File Systems” on page 109.

```
client1 # /usr/bin/kinit root/client1.example.com
Password for root/client1.example.com@EXAMPLE.COM: <Type password>
```

To use the keytab file option, add the `root` principal to the client’s keytab by using `kadmin`:

```
client1 # /usr/sbin/kadmin -p kws/admin
Enter password: <Type kws/admin password>
kadmin: ktadd root/client1.example.com
kadmin: Entry for principal root/client.example.com with
kvno 3, encryption type DES-CBC-CRC added to keytab
WRFILE:/etc/krb5/krb5.keytab
kadmin: quit
```

**8. If you want the client to warn users about Kerberos ticket expiration, create an entry in the `/etc/krb5/warn.conf` file.**

See the `warn.conf(4)` man page for more information.

## Example-Setting Up a SEAM Client Using a Non-SEAM KDC

It is possible to set up a SEAM client to work with a non-SEAM KDC. In this case, a line must be included in the `/etc/krb5/krb5.conf` file in the `realms` section. This line changes the protocol that is used when the client is communicating with the Kerberos password-changing server. The format of this line follows.

```
[realms]
    EXAMPLE.COM = {
        kdc = kdc1.example.com
        kdc = kdc2.example.com
        admin_server = kdc1.example.com
        kpasswd_protocol = SET_CHANGE
    }
```

## Setting Up Root Authentication to Mount NFS File Systems

If users want to access a non-Kerberized NFS file system, either the NFS file system can be mounted as `root`, or the file system can be accessed automatically through the automounter whenever users access it (without requiring `root` permissions).

Mounting a Kerberized NFS file system is very much the same, but it does incur an additional obstacle. To mount a Kerberized NFS file system, users must use the `kinit` command as `root` to obtain credentials for the client's `root` principal, because a client's `root` principal is typically not in the client's keytab. This step is required even when the automounter is set up. This step also forces all users to know their system's `root` password and the `root` principal's password.

To bypass this step, you can add a client's `root` principal to the client's keytab file, which automatically provides credentials for `root`. Although this solution enables users to mount NFS file systems without running the `kinit` command and enhances ease-of-use, it is a security risk. For example, if someone gains access to a system with the `root` principal in its keytab, this person can obtain credentials for `root`. So make sure that you take the appropriate security precautions. See “Administering Keytab Files” on page 166 for more information.

---

## Synchronizing Clocks between KDCs and SEAM Clients

All hosts that participate in the Kerberos authentication system must have their internal clocks synchronized within a specified maximum amount of time (known as *clock skew*). This requirement provides another Kerberos security check. If the clock skew is exceeded between any of the participating hosts, client requests are rejected.

The clock skew also determines how long application servers must keep track of all Kerberos protocol messages, in order to recognize and reject replayed requests. So, the longer the clock skew value, the more information that application servers have to collect.

The default value for the maximum clock skew is 300 seconds (five minutes). You can change this default in the `libdefaults` section of the `krb5.conf` file.

---

**Note** – For security reasons, do not increase the clock skew beyond 300 seconds.

---

Since it is important to maintain synchronized clocks between the KDCs and SEAM clients, you should use the Network Time Protocol (NTP) software to synchronize them. NTP public domain software from the University of Delaware is included in the Solaris software, starting with the Solaris 2.6 release.

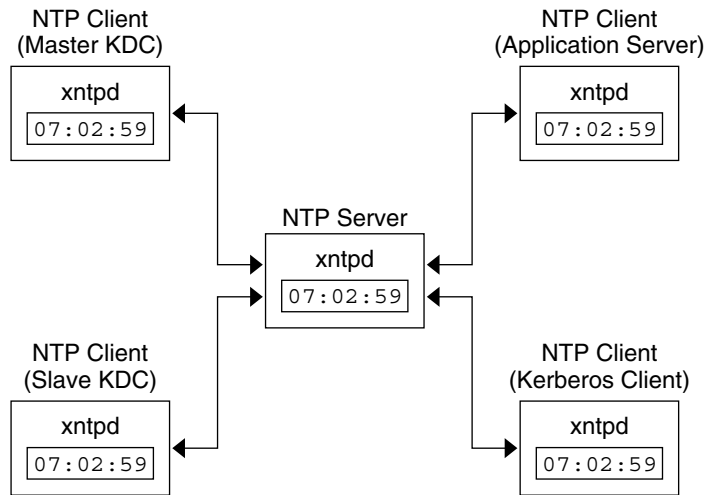
---

**Note** – Another way to synchronize clocks is to use the `rdate` command and `cron` jobs, a process that can be less involved than using NTP. However, this section will continue to focus on using NTP. And, if you use the network to synchronize the clocks, the clock synchronization protocol must itself be secure.

---

NTP enables you to manage precise time or network clock synchronization, or both, in a network environment. NTP is basically a server/client implementation. You pick one system to be the master clock (the NTP server). Then, you set up all your other systems (the NTP clients) to synchronize their clocks with the master clock.

To synchronize the clocks, NTP uses the `xntpd` daemon, which sets and maintains a UNIX system time-of-day in agreement with Internet standard time servers. The following shows an example of this server/client NTP implementation.



**FIGURE 8-1** Synchronizing Clocks by Using NTP

To ensure that the KDCs and SEAM clients maintain synchronized clocks, implement the following steps:

1. Set up an NTP server on your network (this server can be any system, except the master KDC). See “Managing Network Time Protocol (Tasks)” in *System Administration Guide: Resource Management and Network Services* to find the NTP server task.

2. As you configure the KDCs and SEAM clients on the network, set them up to be NTP clients of the NTP server. See “Managing Network Time Protocol (Tasks)” in *System Administration Guide: Resource Management and Network Services* to find the NTP client task.

---

## Swapping a Master KDC and a Slave KDC

You should use the procedures in this section to make the swap of a master KDC with a slave KDC easier. You should swap the master KDC with a slave KDC only if the master KDC server fails for some reason, or if the master KDC needs to be re-installed (for example, because new hardware is installed).

### ▼ How to Configure a Swappable Slave KDC

Perform this procedure on the slave KDC server that you want to have available to become the master KDC.

1. **Use alias names for the master KDC and the swappable slave KDC during the KDC installation.**

When you define the host names for the KDCs, make sure that each system has an alias included in DNS. Also, use the alias names when you define the hosts in the `/etc/krb5/krb5.conf` file.

2. **Follow the steps to install a slave KDC.**

Prior to any swap, this server should function as any other slave KDC in the realm. See “How to Configure a Slave KDC” on page 95 for instructions.

3. **Move the master KDC commands.**

To prevent the master KDC commands from being run from this slave KDC, move the `kprop`, `kadmind` and `kadmin.local` commands to a reserved place.

```
kdc4 # mv /usr/lib/krb5/kprop /usr/lib/krb5/kprop.save
kdc4 # mv /usr/lib/krb5/kadmind /usr/lib/krb5/kadmind.save
kdc4 # mv /usr/sbin/kadmin.local /usr/sbin/kadmin.local.save
```

4. **Comment out the `kprop` line in the root crontab file.**

This step prevents the slave KDC from propagating its copy of the KDC database.

```
kdc4 # crontab -e
#ident "@(#)root 1.19 98/07/06 SMI" /* SVr4.0 1.1.3.1 */
#
```

```

# The root crontab should be used to perform accounting data collection.
#
# The rtc command is run to adjust the real time clock if and when
# daylight savings time changes.
#
10 3 * * 0,4 /etc/cron.d/logchecker
10 3 * * 0 /usr/lib/newsyslog
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [ -x /usr/sbin/rtc ] && /usr/sbin/rtc -c > /dev/null 2>&1
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
#10 3 * * * /usr/lib/krb5kprop_script kdc1.example.sun.com #SUNWkr5ma

```

## ▼ How to Swap a Master KDC and a Slave KDC

This procedure requires that the slave KDC server has been set up as a swappable slave (see “How to Configure a Swappable Slave KDC” on page 111). In this procedure, the master KDC server that is being swapped out is named `kdc1`. The slave KDC that will become the new master KDC is named `kdc4`.

1. On the old master KDC, kill the `kadmind` process.

```
kdc1 # /etc/init.d/kdc.master stop
```

When you kill the `kadmind` process, you prevent any changes from being made to the KDC database.

2. On the old master KDC, comment out the `kprop` line in the root crontab file.

```

kdc1 # crontab -e
#ident  "@(#)root      1.19    98/07/06 SMI"    /* SVr4.0 1.1.3.1    */
#
# The root crontab should be used to perform accounting data collection.
#
# The rtc command is run to adjust the real time clock if and when
# daylight savings time changes.
#
10 3 * * 0,4 /etc/cron.d/logchecker
10 3 * * 0 /usr/lib/newsyslog
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [ -x /usr/sbin/rtc ] && /usr/sbin/rtc -c > /dev/null 2>&1
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
#10 3 * * * /usr/lib/krb5/kprop_script kdc2.example.sun.com #SUNWkr5ma

```

This step prevents the old master from propagating its copy of the KDC database.

3. On the old master KDC, run `kprop_script` to back up and propagate the database.

```

kdc1 # /usr/lib/krb5/kprop_script kdc4.example.com
Database propagation to kdc4.example.com: SUCCEEDED

```

4. On the old master KDC, move the master KDC commands.



To prevent the master KDC commands from being run, move the `kprop`, `kadmind` and `kadmin.local` commands to a reserved place.

```
kdc4 # mv /usr/lib/krb5/kprop /usr/lib/krb5/kprop.save
kdc4 # mv /usr/lib/krb5/kadmind /usr/lib/krb5/kadmind.save
kdc4 # mv /usr/sbin/kadmin.local /usr/sbin/kadmin.local.save
kdc4 # mv /etc/krb5/kadm5.ac1 /etc/krb5/kadm5.ac1.save
```

**5. On the DNS server, change the alias names for the master KDC.**

To change the servers, edit the `example.com` zone file and change the entry for `masterkdc`.

```
masterkdc IN CNAME kdc4
```

**6. On the DNS server, restart the Internet domain name server.**

Run the following command on both servers to get the new alias information:

```
# pkill -1 in.named
```

**7. On the new master KDC, move the master KDC commands.**

```
kdc4 # mv /usr/lib/krb5/kprop.save /usr/lib/krb5/kprop
kdc4 # mv /usr/lib/krb5/kadmind.save /usr/lib/krb5/kadmind
kdc4 # mv /usr/sbin/kadmin.local.save /usr/sbin/kadmin.local
```

**8. On the new master KDC, edit the Kerberos access control list file (`kadm5.ac1`).**

Once populated, the `/etc/krb5/kadm5.ac1` file should contain all principal names that are allowed to administer the KDC. The first entry that is added might look similar to the following:

```
kws/admin@EXAMPLE.COM *
```

This entry gives the `kws/admin` principal in the `EXAMPLE.COM` realm the ability to modify principals or policies in the KDC. The default installation includes an asterisk (\*) to match all admin principals. This default could be a security risk, so it is more secure to include a list of all of the admin principals. See the `kadm5.ac1(4)` man page for more information.

**9. On the new master KDC, create a keytab file for `kadmin` by using `kadmin.local`.**

This command sequence creates a special keytab file with principal entries for `admin` and `changepw`. These principals are needed for the `kadmind` service.

```
kdc4 # /usr/sbin/kadmin.local
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab kadmin/kdc4.example.com
Entry for principal kadmin/kdc4.example.com with kvno 3, encryption type DES-CBC-CRC
added to keytab WRFILE:/etc/krb5/kadm5.keytab.
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab changepw/kdc4.example.com
Entry for principal changepw/kdc4.example.com with kvno 3, encryption type DES-CBC-CRC
added to keytab WRFILE:/etc/krb5/kadm5.keytab.
kadmin.local: quit
```

**10. On the new master KDC, start `kadmind`.**

```
kdc4 # /etc/init.d/kdc.master start
```

## 11. Enable the kprop line in the root crontab file.

```
kdc4 # crontab -e
#ident  "@(#)root      1.19    98/07/06 SMI"    /* SVr4.0 1.1.3.1    */
#
# The root crontab should be used to perform accounting data collection.
#
# The rtc command is run to adjust the real time clock if and when
# daylight savings time changes.
#
10 3 * * 0,4 /etc/cron.d/logchecker
10 3 * * 0 /usr/lib/newsyslog
15 3 * * 0 /usr/lib/fs/nfs/nfsfind
1 2 * * * [ -x /usr/sbin/rtc ] && /usr/sbin/rtc -c > /dev/null 2>&1
30 3 * * * [ -x /usr/lib/gss/gsscred_clean ] && /usr/lib/gss/gsscred_clean
10 3 * * * /usr/lib/krb5/kprop_script kdc1.example.sun.com #SUNWkr5ma
```

---

# Administering the Kerberos Database

The Kerberos database is the backbone of Kerberos and must be maintained properly. This section provides some procedures on how to administer the Kerberos database, such as backing up and restoring the database, setting up parallel propagation, and administering the stash file. The steps to initially set up the database are in “How to Configure a Master KDC” on page 91.

## Backing Up and Propagating the Kerberos Database

Propagating the Kerberos database from the master KDC to the slave KDCs is one of the most important configuration tasks. If propagation doesn't happen often enough, the master KDC and the slave KDCs will lose synchronization. So, if the master KDC goes down, the slave KDCs will not have the most recent database information. Also, if a slave KDC has been configured as a master KDC for purposes of load balancing, the clients that use that slave KDC as a master KDC will not have the latest information. Therefore, you must make sure that propagation occurs often enough, based on how often you change the Kerberos database.

When you configure the master KDC, you set up the `kprop_script` command in a cron job to automatically back up the Kerberos database to the `/var/krb5/slave_datatrans` dump file and propagate it to the slave KDCs. But, as with any file, the Kerberos database can become corrupted. If data corruption occurs on a slave KDC, you might never notice, since the next automatic propagation

of the database installs a fresh copy. However, if corruption occurs on the master KDC, the corrupted database is propagated to all of the slave KDCs during the next propagation. And, the corrupted backup overwrites the previous uncorrupted backup file on the master KDC.

Because there is no “safe” backup copy in this scenario, you should also set up a `cron` job to periodically copy the `slave_datatrans` dump file to another location or to create another separate backup copy by using the `dump` command of `kdb5_util`. Then, if your database becomes corrupted, you can restore the most recent backup on the master KDC by using the `load` command of `kdb5_util`.

Another important note: Because the database dump file contains principal keys, you need to protect the file from being accessed by unauthorized users. By default, the database dump file has read and write permissions only as `root`. To protect against unauthorized access, use only the `kprop` command to propagate the database dump file, which encrypts the data that is being transferred. Also, `kprop` propagates the data only to the slave KDCs, which minimizes the chance of accidentally sending the database dump to unauthorized hosts.



---

**Caution** – If the Kerberos database is updated after it has been propagated and if the database subsequently is corrupted before the next propagation, the KDC slaves will not contain the updates. The updates will be lost. For this reason, if you add significant updates to the Kerberos database before a regularly scheduled propagation, you should manually propagate the database to avoid data loss.

---

## The `kpropd.ac1` File

The `kpropd.ac1` file on a KDC provides a list of host principal names, one per line, that specifies the systems from which the KDC can receive an updated database through propagation. If the master KDC is used to propagate all the slave KDCs, the `kpropd.ac1` file on each slave needs to contain only the host principal name of the master KDC.

However, the SEAM installation and subsequent configuration steps in this book instruct you to add the same `kpropd.ac1` file to the master KDC and the slave KDCs. This file contains all the KDC host principal names. This configuration allows you to propagate from any KDC, in case the propagating KDCs become temporarily unavailable. And, by keeping an identical copy on all KDCs, you make the configuration easy to maintain.

## The kprop\_script Command

The `kprop_script` command uses the `kprop` command to propagate the Kerberos database to other KDCs. If the `kprop_script` command is run on a slave KDC, it propagates the slave KDC's copy of the Kerberos database to other KDCs. The `kprop_script` accepts a list of host names for arguments, separated by spaces, which denote the KDCs to propagate.

When the `kprop_script` is run, it creates a backup of the Kerberos database to the `/var/krb5/slave_datatrans` file and copies the file to the specified KDCs. The Kerberos database is locked until the propagation is finished.

## ▼ How to Back Up the Kerberos Database

1. Become superuser on the master KDC.
2. Back up the Kerberos database by using the `dump` command of the `kdb5_util` command.

```
# /usr/sbin/kdb5_util dump [-verbose] [-d dbname] [filename [principals...]]
```

<code>-verbose</code>	Prints the name of each principal and policy that is being backed up.
<code>dbname</code>	Defines the name of the database to back up. Note that “.db” is appended to whatever database name is specified, and you can specify an absolute path for the file. If the <code>-d</code> option is not specified, the default database name is <code>/var/krb5/principal</code> , which actually becomes <code>/var/krb5/principal.db</code> .
<code>filename</code>	Defines the file that is used to back up the database. You can specify an absolute path for the file. If you don't specify a file, the database is dumped to standard output.
<code>principal</code>	Defines a list of one or more principals (separated by a space) to back up. You must use fully-qualified principal names. If you don't specify any principals, the entire database is backed up.

## Example—Backing Up the Kerberos Database

In the following example, the Kerberos database is backed up to a file called `dumpfile`. Because the `-verbose` option is specified, each principal is printed as it is backed up.

```
# kdb5_util dump -verbose dumpfile
kadmin/kdc1.eng.example.com@ENG.EXAMPLE.COM
krbtgt/eng.example.com@ENG.EXAMPLE.COM
kadmin/history@ENG.EXAMPLE.COM
```

```
pak/admin@ENG.EXAMPLE.COM
pak@ENG.EXAMPLE.COM
changepw/kdc1.eng.example.com@ENG.EXAMPLE.COM
```

In the following example, the pak and pak/admin principals are backed up from the Kerberos database.

```
# kdb5_util dump -verbose dumpfile pak/admin@ENG.EXAMPLE.COM pak@ENG.EXAMPLE.COM
pak/admin@ENG.EXAMPLE.COM
pak@ENG.EXAMPLE.COM
```

## ▼ How to Restore the Kerberos Database

1. Become superuser on the master KDC.
2. Restore the Kerberos database by using the `load` command of `kdb_util`.

```
# /usr/sbin/kdb5_util load [-verbose] [-d dbname] [-update] [filename]
```

<code>-verbose</code>	Prints the name of each principal and policy that is being restored.
<code>dbname</code>	Defines the name of the database to restore. Note that “.db” is appended to whatever database name is specified, and you can specify an absolute path for the file. If the <code>-d</code> option is not specified, the default database name is <code>/var/krb5/principal</code> , which actually becomes <code>/var/krb5/principal.db</code> .
<code>-update</code>	Updates the existing database. Otherwise, a new database is created or the existing database is overwritten.
<code>filename</code>	Defines the file from which to restore the database. You can specify an absolute path for the file.

### Example—Restoring the Kerberos Database

In the following example, the database called `database1.db` is restored into the current directory from the `dumpfile` file. Since the `-update` option isn't specified, a new database is created by the restore.

```
# kdb5_util load -d database1 dumpfile
```

## ▼ How to Manually Propagate the Kerberos Database to the Slave KDCs

This procedure shows you how to propagate the Kerberos database by using the `kprop` command. Use this procedure if you need to synchronize a slave KDC with the master KDC outside the periodic `cron` job. And, unlike the `kprop_script`, you can use `kprop` to propagate just the current database backup without first making a new backup of the Kerberos database.

1. **Become superuser on the master KDC.**
2. **(Optional) Back up the database by using the `kdb5_util` command.**
3. **Propagate the database to a slave KDC by using the `kprop` command.**

```
# /usr/sbin/kdb5_util dump /var/krb5/slave_datatrans
```

```
# /usr/lib/krb5/kprop -f /var/krb5/slave_datatrans slave_KDC
```

If you want to back up the database and propagate it to a slave KDC outside the periodic `cron` job, you can also use the `kprop_script` command as follows:


```
# /usr/lib/krb5/kprop_script slave_KDC
```

## Setting Up Parallel Propagation

In most cases, the master KDC is used exclusively to propagate its Kerberos database to the slave KDCs. However, if your site has many slave KDCs, you might consider load-sharing the propagation process, known as *parallel propagation*.

Parallel propagation allows specific slave KDCs to share the propagation duties with the master KDC. This sharing of duties enables the propagation to be done faster and to lighten the work for the master KDC.

For example, say your site has one master KDC and six slave KDCs (shown in Figure 8-2), where `slave-1` through `slave-3` consist of one logical grouping and `slave-4` through `slave-6` consist of another logical grouping. To set up parallel propagation, you could have the master KDC propagate the database to `slave-1` and `slave-4`, and those KDC slaves could in turn propagate the database to the KDC slaves in their group.



**FIGURE 8-2** Example of Parallel Propagation Configuration

## ▼ How to Set Up Parallel Propagation

The following is not a detailed step-by-step procedure, but a high-level list of configuration steps to enable parallel propagation.

1. On the master KDC, change the `kprop_script` entry in its cron job to include arguments for only the KDC slaves that will perform the succeeding propagation (*propagation slaves*).
2. On each propagation slave, add a `kprop_script` entry to its cron job, which must include arguments for the slaves to propagate. To successfully propagate in parallel, the cron job should be set up to run after the propagation slave is itself propagated with the new Kerberos database.

---

**Note** – How long it will take for a propagation slave to be propagated depends on factors such as network bandwidth and the size of the database.

---

3. On each slave KDC, set up the appropriate permissions to be propagated. This step is done by adding the host principal name of its propagating KDC to its `kpropd.acl` file.

### Example—Setting Up Parallel Propagation

Using the example in Figure 8-2, the master KDC's `kprop_script` entry would look similar to the following:

```
0 3 * * * /usr/lib/krb5/kprop_script slave-1.example.com slave-4.example.com
```

The slave-1's `kprop_script` entry would look similar to the following:

```
0 4 * * * /usr/lib/krb5/kprop_script slave-2.example.com slave-3.example.com
```

Note that the propagation on the slave starts an hour after it is propagated by the master.

The `kpropd.acl` file on the propagation slaves would contain the following entry:

```
host/master.example.com@EXAMPLE.COM
```

The `kpropd.acl` file on the KDC slaves being propagated by `slave-1` would contain the following entry:

```
host/slave-1.example.com@EXAMPLE.COM
```

## Administering the Stash File

The stash file contains the master key for the Kerberos database, which is automatically created when you create a Kerberos database. If the stash file gets corrupted, you can use the `stash` command of the `kdb5_util` utility to replace the corrupted file. The only time you should need to remove a stash file is after removing the Kerberos database with the `destroy` command of `kdb5_util`. Because the stash file isn't automatically removed with the database, you have to remove it to finish the cleanup.

### ▼ How to Remove a Stash File

1. **Become superuser on the KDC that contains the stash file.**
2. **Remove the stash file.**

```
# rm stash-file
```

In this example, *stash-file* is the path to the stash file. By default, the stash file is located at `/var/krb5/.k5.realm`.

If you need to re-create the stash file, you can use the `-f` option of the `kdb5_util` command.

---

## Increasing Security

Follow these steps to increase security on SEAM application servers and on KDC servers.

### ▼ How to Restrict Access to KDC Servers

Both master KDC servers and slave KDC servers have copies of the KDC database stored locally. Restricting access to these servers so that the databases are secure is important to the overall security of the SEAM installation.



### 1. Disable remote services in the `/etc/inetd.conf` file.

To provide a secure KDC server, all nonessential network services should be disabled by commenting out the entry that starts the service in the `/etc/inetd.conf` file. In most circumstances, the only services that would need to run would be `time` and `krdb5_kprop`. In addition, any services that use loopback tli (`ticlts`, `ticotsord`, and `ticots`) can be left enabled. After you edit the file, it should look similar to the following (to shorten the example many comments have been removed):

```
kdc1 # cat /etc/inetd.conf
#
#ident  "@(#)inetd.conf 1.33      98/06/02 SMI"      /* SVr4.0 1.5  */
.
.
#name    dgram    udp      wait     root     /usr/sbin/in.tnamed    in.tnamed
#
#shell   stream   tcp      nowait   root     /usr/sbin/in.rshd      in.rshd
#login   stream   tcp      nowait   root     /usr/sbin/in.rlogind   in.rlogind
#exec    stream   tcp      nowait   root     /usr/sbin/in.rexecd    in.rexecd
#comsat  dgram    udp      wait     root     /usr/sbin/in.comsat    in.comsat
#talk    dgram    udp      wait     root     /usr/sbin/in.talkd     in.talkd
#
#uucp    stream   tcp      nowait   root     /usr/sbin/in.uucpd     in.uucpd
#
#finger  stream   tcp      nowait   nobody   /usr/sbin/in.fingerd   in.fingerd
#
# Time service is used for clock synchronization.
#
time     stream   tcp      nowait   root     internal
time     dgram    udp      wait     root     internal
#
.
.
#
100234/1 tli  rpc/ticotsord wait   root     /usr/lib/gss/gssd      gssd
#dtspc   stream   tcp      nowait   root     /usr/dt/bin/dtspcd     /usr/dt/bin/dtspcd
#100068/2-5 dgram  rpc/udp  wait     root     /usr/dt/bin/rpc.cmsd   rpc.cmsd
100134/1 tli  rpc/ticotsord wait   root     /usr/lib/ktkt_warnd    kwarnd
krb5_prop stream  tcp      nowait   root     /usr/lib/krb5/kpropd   kpropd
```

Reboot the KDC server after the changes are made.

### 2. Restrict access to the hardware that supports the KDC.

In order to restrict physical access, make sure that the KDC server and its monitor are located in a secure facility. Users should not be able to access this server in any way.

### 3. Store KDC database backups on local disks or on the KDC slaves.

Make tape backups of your KDC only if the tapes are stored securely. Follow the same practice for copies of keytab files. It would be best to store these files on a local file system that is not shared to other systems. The storage file system can be on either the master KDC server or any of the slave KDCs.



# SEAM Error Messages and Troubleshooting

---

This chapter provides resolutions for error messages that you might receive when you use SEAM, as well as some troubleshooting tips for various problems. This is a list of the error message and troubleshooting information in this chapter.

- “SEAM Administration Tool Error Messages” on page 123
- “Common SEAM Error Messages (A-M)” on page 124
- “Common SEAM Error Messages (N-Z)” on page 130
- “Problems Mounting a Kerberized NFS File System” on page 133
- “Problems Authenticating as root” on page 134

---

## SEAM Error Messages

This section provides information about SEAM error messages, including why each error occurs and a way to fix it.

### SEAM Administration Tool Error Messages

Unable to view the list of principals or policies; use the Name field.

**Cause:** The admin principal that you logged in with does not have the list privilege (1) in the Kerberos ACL file (`kadm5.ac1`), so you cannot view the principal list or policy list.

**Solution:** You must enter the principal and policy names in the Name field to work on them, or you need to log on with a principal that has the appropriate privileges.

JNI: Java array creation failed  
JNI: Java class lookup failed  
JNI: Java field lookup failed  
JNI: Java method lookup failed  
JNI: Java object lookup failed  
JNI: Java object field lookup failed  
JNI: Java string access failed  
JNI: Java string creation failed

**Cause:** A serious problem exists with the Java Native Interface that is used by the SEAM Administration Tool (gkadmin).

**Solution:** Exit gkadmin and restart it. If the problem persists, please report a bug.

## Common SEAM Error Messages (A-M)

This section provides an alphabetical list (A-M) of common error messages for the SEAM commands, SEAM daemons, PAM framework, GSS interface, the NFS service, and the Kerberos library.

*major\_error minor\_error gssapi error importing name*

**Cause:** An error occurred while a service name was being imported.

**Solution:** Make sure that the service principal is in the host's keytab file.

Bad krb5 admin server hostname while initializing kadmind interface

**Cause:** An invalid host name is configured for `admin_server` in the `krb5.conf` file.

**Solution:** Make sure that the correct host name for the master KDC is specified on the `admin_server` line in the `krb5.conf` file.

Cannot contact any KDC for requested realm

**Cause:** No KDC responded in the requested realm.

**Solution:** Make sure that at least one KDC (either the master or slave) is reachable or that the `krb5kdc` daemon is running on the KDCs. Check the `/etc/krb5/krb5.conf` file for the list of configured KDCs (`kdc = kdc_name`).

Cannot determine realm for host

**Cause:** Kerberos cannot determine the realm name for the host.

**Solution:** Make sure that there is a default realm name, or that the domain name mappings are set up in the Kerberos configuration file (`krb5.conf`).

Cannot find KDC for requested realm

**Cause:** No KDC was found in the requested realm.

**Solution:** Make sure that the Kerberos configuration file (`krb5.conf`) specifies a KDC in the realm section.

cannot initialize realm *realm\_name*

**Cause:** The KDC might not have a stash file.

**Solution:** Make sure that the KDC has a stash file. If not, create a stash file by using the `kdb5_util` command, and try running the `krb5kdc` command again. The easiest way to start `krb5kdc` is to run the `/etc/init.d/kdc` script.

Cannot resolve KDC for requested realm

**Cause:** Kerberos cannot determine any KDC for the realm.

**Solution:** Make sure that the Kerberos configuration file (`krb5.conf`) specifies a KDC in the realm section.

Cannot reuse password

**Cause:** The password that you entered has been used before by this principal.

**Solution:** Choose a password that has not been chosen before, at least not within the number of passwords that are kept in the KDC database for each principal (this policy is enforced by the principal's policy).

Can't get forwarded credentials

**Cause:** Credential forwarding could not be established.

**Solution:** Make sure that the principal has forwardable credentials.

Can't open/find Kerberos configuration file

**Cause:** The Kerberos configuration file (`krb5.conf`) was unavailable.

**Solution:** Make sure that the `krb5.conf` file is available in the correct location and has the correct permissions. This file should be writable by `root` and readable by everyone else.

Client/server realm mismatch in initial ticket request

**Cause:** A realm mismatch between the client and server occurred in the initial ticket request.

**Solution:** Make sure that the server you are communicating with is in the same realm as the client, or that the realm configurations are correct.

Client or server has a null key

**Cause:** The principal has a null key.

**Solution:** Modify the principal to have a non-null key by using the `cpw` command of `kadmin`.

Communication failure with server while initializing `kadmin` interface

**Cause:** The host that was entered for the admin server, also called the master KDC, did not have the `kadmind` daemon running.

**Solution:** Make sure that you specified the correct host name for the master KDC. If you specified the correct host name, make sure that `kadmind` is running on the master KDC that you specified.

Credentials cache file permissions incorrect

**Cause:** You do not have the appropriate read or write permissions on the credentials cache (`/tmp/krb5cc_uid`).

**Solution:** Make sure that you have read and write permissions on the credentials cache.

Credentials cache I/O operation failed XXX

**Cause:** Kerberos had a problem writing to the system's credentials cache (`/tmp/krb5cc_uid`).

**Solution:** Make sure that the credentials cache has not been removed, and that there is space left on the device by using the `df` command.

Decrypt integrity check failed

**Cause:** You might have an invalid ticket.

**Solution:**

1. Make sure that your credentials are valid. Destroy your tickets with `kdestroy` and create new tickets with `kinit`.
2. Make sure that the target host has a keytab file with the correct version of the service key. Use `kadmin` to view the key version number of the service principal (for example, `host/FQDN_hostname`) in the Kerberos database. Also, use `klist -k` on the target host to make sure that it has the same key version number.

`df: cannot statvfs filesystem: Invalid argument`

**Cause:** The `df` command cannot access the Kerberized NFS file system, which is currently mounted, to generate its report, because you no longer have the appropriate root credentials. Destroying your credentials for a mounted Kerberized file system does not automatically unmount the file system.

**Solution:** You must create new root credentials to access the Kerberized file system. If you no longer require access to the Kerberized file system, unmount the file system.

failed to obtain credentials cache

**Cause:** During `kadmin` initialization, a failure occurred when `kadmin` tried to obtain credentials for the `admin` principal.

**Solution:** Make sure that you used the correct principal and password when you executed `kadmin`.

Field is too long for this implementation

**Cause:** The message size that was being sent by a Kerberized application was too long. The maximum message size that can be handled by Kerberos is 65535 bytes. In addition, there are limits on individual fields within a protocol message that is sent by Kerberos.

**Solution:** Make sure that your Kerberized applications are sending valid message sizes.

GSS-API (or Kerberos) error

**Cause:** This message is a generic GSS-API or Kerberos error message and can be caused by several different problems.

**Solution:** Check the `/etc/krb5/kdc.log` file to find the more specific GSS-API error message that was logged when this error occurred.

Hostname cannot be canonicalized

**Cause:** Kerberos cannot make the host name fully qualified.

**Solution:** Make sure that the host name is defined in DNS and that the `host-name-to-address` and `address-to-host-name` mappings are consistent.

Illegal cross-realm ticket

**Cause:** The ticket sent did not have the correct cross-realms. The realms might not have the correct trust relationships set up.

**Solution:** Make sure that the realms you are using have the correct trust relationships.

Improper format of Kerberos configuration file

**Cause:** The Kerberos configuration file (`krb5.conf`) has invalid entries.

**Solution:** Make sure that all the relations in the `krb5.conf` file are followed by the "=" sign and a value. Also, verify that the brackets are present in pairs for each subsection.

Inappropriate type of checksum in message

**Cause:** The message contained an invalid checksum type.

**Solution:** Check which valid checksum types are specified in the `krb5.conf` and `kdc.conf` files.

Incorrect net address

**Cause:** There was a mismatch in the network address. The network address in the ticket that was being forwarded was different from the network address where the ticket was processed. This message might occur when tickets are being forwarded.

**Solution:** Make sure that the network addresses are correct. Destroy your tickets with `kdestroy`, and create new tickets with `kinit`.

Invalid flag for file lock mode

**Cause:** An internal Kerberos error occurred.

**Solution:** Please report a bug.

Invalid message type specified for encoding

**Cause:** Kerberos could not recognize the message type that was sent by the Kerberized application.

**Solution:** If you are using a Kerberized application that was developed by your site or a vendor, make sure that it is using Kerberos correctly.

Invalid number of character classes

**Cause:** The password that you entered for the principal does not contain enough password classes, as enforced by the principal's policy.

**Solution:** Make sure that you enter a password with the minimum number of password classes that the policy requires.

KADM err: Memory allocation failure

**Cause:** There is not enough memory to run `kadmin`.

**Solution:** Free up memory and try running `kadmin` again.

KDC can't fulfill requested option

**Cause:** The KDC did not allow the requested option. A possible problem might be that postdating or forwardable options were being requested, and the KDC did not allow it. Another problem might be that you requested the renewal of a TGT, but you didn't have a renewable TGT.

**Solution:** Determine if you are requesting an option that either the KDC does not allow or if you are requesting a type of ticket that is not available.

KDC policy rejects request

**Cause:** The KDC policy did not allow the request. For example, the request to the KDC did not have an IP address in its request, or forwarding was requested, but the KDC did not allow it.



**Solution:** Make sure that you are using `kinit` with the correct options. If necessary, modify the policy that is associated with the principal or change the principal's attributes to allow the request. You can modify the policy or principal by using `kadmin`.

KDC reply did not match expectations

**Cause:** The KDC reply did not contain the expected principal name, or other values in the response were incorrect.

**Solution:** Make sure that the KDC you are communicating with complies with RFC1510, or that the request you are sending is a Kerberos V5 request, or that the KDC is available.

Key table entry not found

**Cause:** There is no entry for the service principal in the network application server's keytab file.

**Solution:** Add the appropriate service principal to the server's keytab file so that it can provide the Kerberized service.

Key version number for principal in key table is incorrect

**Cause:** A principal's key version is different in the keytab file and in the Kerberos database. Either a service's key has been changed, or you might be using an old service ticket.

**Solution:** If a service's key has been changed (for example, by using `kadmin`), you need to extract the new key and store it in the host's keytab file where the service is running.

Alternately, you might be using an old service ticket that has an older key. You might want to run the `kdestroy` command and then the `kinit` command again.

login: load\_modules: can not open module  
/usr/lib/security/pam\_krb5.so.1

**Cause:** Either the Kerberos PAM module is missing or it is not a valid executable binary.

**Solution:** Make sure that the Kerberos PAM module is in the `/usr/lib/security` directory and that it is a valid executable binary. Also, make sure that the `/etc/pam.conf` file contains the correct path to `pam_krb5.so.1`.

Looping detected inside `krb5_get_in_tkt`

**Cause:** Kerberos made several attempts to get the initial tickets but failed.

**Solution:** Make sure that at least one KDC is responding to authentication requests.

Master key does not match database

**Cause:** The loaded database dump was not created from a database that contains the master key, which is located in `/var/krb5/.k5.REALM`.

**Solution:** Make sure that the master key in the loaded database dump matches the master key that is located in `/var/krb5/.k5.REALM`.

Matching credential not found

**Cause:** The matching credential for your request was not found. Your request requires credentials that are unavailable in the credentials cache.

**Solution:** Destroy your tickets with `kdestroy` and create new tickets with `kinit`.

Message out of order

**Cause:** Messages that were sent using sequential-order privacy arrived out of order. Some messages might have been lost in transit.

**Solution:** You should reinitialize the Kerberos session.

Message stream modified

**Cause:** There was a mismatch between the computed checksum and the message checksum. The message might have been modified while in transit, which can indicate a security leak.

**Solution:** Make sure that the messages are being sent across the network correctly. Since this message can also indicate the possible tampering of messages while they are being sent, destroy your tickets using `kdestroy` and reinitialize the Kerberos services that you are using.

## Common SEAM Error Messages (N-Z)

This section provides an alphabetical list (N-Z) of common error messages for the SEAM commands, SEAM daemons, PAM framework, GSS interface, the NFS service, and the Kerberos library.

No credentials cache file found

**Cause:** Kerberos could not find the credentials cache (`/tmp/krb5cc_uid`).

**Solution:** Make sure that the credential file exists and is readable. If it isn't, try performing the `kinit` again.

Operation requires "*privilege*" privilege

**Cause:** The admin principal that was being used does not have the appropriate privilege configured in the `kadm5.acl` file.

**Solution:** Use a principal that has the appropriate privileges. Or, configure the principal that was being used to have the appropriate privileges by modifying the `kadm5 .acl` file. Usually, a principal with `"/admin"` as part of its name has the appropriate privileges.

PAM-KRB5: Kerberos V5 authentication failed: password incorrect

**Cause:** Your UNIX password and Kerberos passwords are different. Most non-Kerberized commands, such as `login`, are set up through PAM to automatically authenticate with Kerberos by using the same password that you specified for your UNIX password. If your passwords are different, the Kerberos authentication fails.

**Solution:** You must enter your Kerberos password when prompted.

Password is in the password dictionary

**Cause:** The password that you entered is in a password dictionary that is being used. Your password is not a good choice for a password.

**Solution:** Choose a password that has a mix of password classes.

Permission denied in replay cache code

**Cause:** The system's replay cache could not be opened. The server might have been first run under a user ID different than your current user ID.

**Solution:** Make sure that the replay cache has the appropriate permissions. The replay cache is stored on the host where the Kerberized server application is running (`/usr/tmp/rc_service_name`). Instead of changing the permissions on the current replay cache, you can also remove the replay cache before you run the Kerberized server under a different user ID.

Protocol version mismatch

**Cause:** Most likely, a Kerberos V4 request was sent to the KDC. SEAM supports only the Kerberos V5 protocol.

**Solution:** Make sure that your applications are using the Kerberos V5 protocol.

Request is a replay

**Cause:** The request has already been sent to this server and processed. The tickets might have been stolen, and someone else is trying to reuse the tickets.

**Solution:** Wait for a few minutes and reissue the request.

Requested principal and ticket don't match

**Cause:** The service principal that you are connecting to and the service ticket that you have do not match.

**Solution:** Make sure that DNS is functioning properly. If you are using another vendor's software, make sure that the software is using principal names correctly.

Requested protocol version not supported

**Cause:** Most likely, a Kerberos V4 request was sent to the KDC. SEAM supports only the Kerberos V5 protocol.

**Solution:** Make sure that your applications are using the Kerberos V5 protocol.

Required parameters in krb5.conf missing while initializing kadmin interface

**Cause:** There is a missing parameter (such as the `admin_server` parameter) in the `krb5.conf` file.

**Solution:** Determine which parameter is missing and add it to the `krb5.conf` file.

Server rejected authentication (during sendauth exchange)

**Cause:** The server that you are trying to communicate with rejected the authentication. Most often this error occurs during Kerberos database propagation. Some common causes might be problems with the `kpropd.acl` file, DNS, or the keytab file.

**Solution:** If you get this error when you are running applications other than `kprop`, investigate whether the server's keytab file is correct.

Set gss service nfs@<host> failed. Check nfs service credential.

**Cause:** This message is generated by `syslog` after a `share` command has failed with an "invalid argument" message. The most likely cause of this message is that either there is no keytab file or that there is no NFS service principle in the keytab file.

**Solution:** To isolate the problem, run `klist -k` to see if the keytab file exists and if there is an NFS service principal for the host in the keytab file.

The ticket isn't for us

Ticket/authenticator don't match

**Cause:** There was a mismatch between the ticket and authenticator. The principal name in the request might not have matched the service principal's name, because the ticket was being sent with an FQDN name of the principal while the service expected non-FQDN, or vice versa.

**Solution:** If you get this error when you are running applications other than `kprop`, investigate whether the server's keytab file is correct.

Ticket expired

**Cause:** Your ticket times have expired.

**Solution:** Destroy your tickets with `kdestroy` and create new tickets with `kinit`.

Ticket is ineligible for postdating

**Cause:** The principal does not allow its tickets to be postdated.

**Solution:** Modify the principal with `kadmin` to allow postdating.

Ticket not yet valid

**Cause:** The postdated ticket is not valid yet.

**Solution:** Create new tickets with the correct date, or wait until the current tickets are valid.

Truncated input file detected

**Cause:** The database dump file that was being used in the operation is not a complete dump file.

**Solution:** Create the dump file again, or use a different database dump file.

Wrong principal in request

**Cause:** There was an invalid principal name in the ticket. This error might indicate a DNS or FQDN problem.

**Solution:** Make sure that the principal of the service matches the principal in the ticket.

---

## SEAM Troubleshooting

This section provides troubleshooting information for the SEAM software.

### Problems Mounting a Kerberized NFS File System

- If mounting a Kerberized NFS file system fails, make sure that the `/var/tmp/rc_nfs` file exists on the NFS server. If the file system is not owned by `root`, remove it and try the mount again.
- If you have a problem accessing a Kerberized NFS file system, make sure that there is an entry for `gssd` in the `inetd.conf` file on your system and the NFS server.
- If you see either the `invalid argument` or `bad directory` error message when you are trying to access a Kerberized NFS file system, the problem might be that you are not using a fully-qualified DNS name when you are trying to mount the NFS file system. The host that is being mounted is not the same as the host name part of the service principal in the server's keytab file.

This problem might also occur if your server has multiple Ethernet interfaces, and you have set up DNS to use a “name per interface” scheme instead of a “multiple address records per host” scheme. For SEAM, you should set up multiple address records per host as follows<sup>1</sup> :

```
my.host.name.    A      1.2.3.4
                 A      1.2.4.4
                 A      1.2.5.4

my-en0.host.name.    A      1.2.3.4
my-en1.host.name.    A      1.2.4.4
my-en2.host.name.    A      1.2.5.4

4.3.2.1          PTR    my.host.name.
4.4.2.1          PTR    my.host.name.
4.5.2.1          PTR    my.host.name.
```

In this example, the setup allows one reference to the different interfaces and allows a single service principal instead of three service principals in the server’s keytab file.

## Problems Authenticating as root

If authentication fails when you try to become superuser on your system and you have already added the `root` principal to your host’s keytab file, there are two potential problems to check. First, make sure that the `root` principal in the keytab file has a fully-qualified name as its instance. If it does, check the `/etc/resolv.conf` file to make sure that the system is correctly set up as a DNS client.

---

<sup>1</sup> Ken Hornstein, “Kerberos FAQ,” [<http://www.nrl.navy.mil/CCS/people./kenh/kerberos-faq.html>], accessed 11 December 1998.

## Administering Principals and Policies (Tasks)

---

This chapter provides procedures for managing principals and the policies that are associated with them. This chapter also shows how to manage a host's keytab file.

This chapter should be used by anyone who needs to administer principals and policies. Before you use this chapter, you should be familiar with principals and policies, including any planning considerations. Refer to Chapter 6 and Chapter 7, respectively.

This is a list of the information in this chapter.

- "Administering Principals" on page 140
- "Administering Policies" on page 153
- "SEAM Tool Reference" on page 162
- "Administering Keytab Files" on page 166

---

## Ways to Administer Principals and Policies

The Kerberos database on the master KDC contains all of your realm's Kerberos principals, their passwords, policies, and other administrative information. To create and delete principals, and to modify their attributes, you can use the `kadmin` or `gkadmin` commands.

The `kadmin` command provides an interactive command-line interface that enables you to maintain Kerberos principals, policies, and keytab files. There are two versions of the `kadmin` command:

- `kadmin`, which uses Kerberos authentication to operate securely from anywhere on the network
- `kadmin.local`, which must be run directly on the master KDC

Other than `kadmin` using Kerberos to authenticate the user, the capabilities of the two versions are identical. The local version is necessary to enable you to set up enough of the database so that you can use the remote version.

Also, SEAM provides the SEAM Administration Tool, `gkadmin`, which is an interactive graphical user interface (GUI) that provides essentially the same capabilities as the `kadmin` command. See “SEAM Administration Tool” on page 136 for more information.

---

## SEAM Administration Tool

The SEAM Administration Tool is an interactive graphical user interface (GUI) that enables you to maintain Kerberos principals and policies. This tool provides much the same capabilities as the `kadmin` command. However, this tool does not support the management of keytab files. You must use the `kadmin` command to administer keytab files, which is described in “Administering Keytab Files” on page 166.

Similar to the `kadmin` command, the SEAM Tool uses Kerberos authentication and encrypted RPC to operate securely from anywhere on the network. The SEAM Tool enables you to do the following:

- Create new principals that are based on default values or existing principals
- Create new policies that are based on existing policies
- Add comments for principals
- Set up default values for creating new principals
- Log in as another principal without exiting the tool
- Print or save principal lists and policy lists
- View and search principal lists and policy lists

The SEAM Tool also provides context-sensitive help and general online help.

The following task maps provide pointers to the various tasks that you can do with the SEAM Tool:

- “Administering Principals (Task Map)” on page 141
- “Administering Policies (Task Map)” on page 153

Also, go to “SEAM Tool Panel Descriptions” on page 162 for descriptions of all the principal attributes and policy attributes that you can either specify or view in the SEAM Tool.



## Command-Line Equivalents of the SEAM Tool

This section lists the `kadmin` commands that provide the same capabilities as the SEAM Tool. These commands can be used without running an X Window system. Even though most procedures in this chapter use the SEAM Tool, many procedures also provide corresponding examples that use the command-line equivalents.

**TABLE 10-1** Command-Line Equivalents of the SEAM Tool

SEAM Tool Procedure	Equivalent <code>kadmin</code> Command
View the list of principals	<code>list_principals</code> or <code>get_principals</code>
View a principal's attributes	<code>get_principal</code>
Create a new principal	<code>add_principal</code>
Duplicate a principal	No command-line equivalent
Modify a principal	<code>modify_principal</code> or <code>change_password</code>
Delete a principal	<code>delete_principal</code>
Set up defaults for creating new principals	No command-line equivalent
View the list of policies	<code>list_policies</code> or <code>get_policies</code>
View a policy's attributes	<code>get_policy</code>
Create a new policy	<code>add_policy</code>
Duplicate a policy	No command-line equivalent
Modify a policy	<code>modify_policy</code>
Delete a policy	<code>delete_policy</code>

## Files Modified by the SEAM Tool

The only file that the SEAM Tool modifies is the `$HOME/.gkadmin` file. This file contains the default values for creating new principals. You can update this file by choosing Properties from the Edit menu.

## Print and Online Help Features of the SEAM Tool

The SEAM Tool provides both print features and online help features. From the Print menu, you can send the following to a printer or a file:

- List of available principals on the specified master KDC
- List of available policies on the specified master KDC
- The currently selected principal or the loaded principal

- The currently selected policy or the loaded policy

From the Help menu, you can access context-sensitive help and general help. When you choose Context-Sensitive Help from the Help menu, the Context-Sensitive Help window is displayed and the tool is switched to help mode. In help mode, when you click on any fields, labels, or buttons on the window, help on that item is displayed in the Help window. To switch back to the tool's normal mode, click Dismiss in the Help window.

You can also choose Help Contents, which opens an HTML browser that provides pointers to the general overview and task information that is provided in this chapter.

## Working With Large Lists in the SEAM Tool

As your site starts to accumulate a large number of principals and policies, the time it takes the SEAM Tool to load and display the principal and policy lists will become increasingly longer. Thus, your our productivity with the tool will increase. There are several ways to work around this problem.

First, you can completely eliminate the time to load the lists by not having the SEAM Tool load the lists. You can set this option by choosing Properties from the Edit menu, and unchecking the Show Lists field. Of course, when the tool doesn't load the lists, it can't display the lists, and you can no longer use the list panels to select principals or policies. Instead, you must type a principal or policy name in the new Name field that is provided, then select the operation that you want to perform on it. In effect, typing a name is equivalent to selecting an item from the list.

Another way to work with large lists is to cache them. In fact, caching the lists for a limited time is set as the default behavior for the SEAM Tool. The SEAM Tool must still initially load the lists into the cache, but after that, the tool can use the cache rather than retrieve the lists again. This option eliminates the need to keep loading the lists from the server, which is what takes so long.

You can set list caching by choosing Properties from the Edit menu. There are two cache settings. You can choose to cache the list forever, or you can specify a time limit when the tool must reload the lists from the server into the cache.

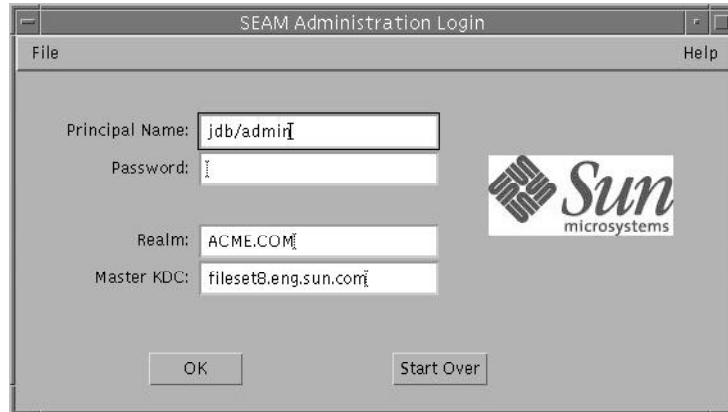
Caching the lists still enables you to use the list panels to select principals and policies, so it doesn't affect how you use the SEAM Tool as the first option does. Also, even though caching doesn't enable you to see the changes of others, you can still see the latest list information based on your changes, since your changes update the lists both on the server and in the cache. And, if you want to update the cache to see other changes and get the lastest copy of the lists, you can use the Refresh menu whenever you want to refresh the cache from the server.

## ▼ How to Start the SEAM Tool

1. Start the SEAM Tool by using the `gkadmin` command.

```
$ /usr/sbin/gkadmin
```

The SEAM Administration Login window is displayed.



2. If you don't want to use the default values, specify new default values.

The window automatically fills in with default values. The default principal name is determined by taking your current identity from the `USER` environment variable and appending `/admin` to it (`username/admin`). The default Realm and Master KDC fields are selected from the `/etc/krb5/krb5.conf` file. If you ever want to retrieve the default values, click `Start Over`.

---

**Note** – The administration operations that each Principal Name can perform are dictated by the Kerberos ACL file, `/etc/krb5/kadm5.acl`. For information about limited privileges, see “Using the SEAM Tool With Limited Kerberos Administration Privileges” on page 165.

---

3. Enter a password for the specified principal name.
4. Click `OK`.

The following window is displayed.



---

## Administering Principals

This section provides the step-by-step instructions to administer principals with the SEAM Tool. This section also provides examples of equivalent command lines, when available.

## Administering Principals (Task Map)

Task	Description	For Instructions
View the list of principals	View the list of principals by clicking the Principals tab.	"How to View the List of Principals" on page 142
View a principal's attributes	View a principal's attributes by selecting the Principal in the Principal List, then clicking the Modify button.	"How to View a Principal's Attributes" on page 144
Create a new principal	Create a new principal by clicking the Create New button in the Principal List panel.	"How to Create a New Principal" on page 146
Duplicate a principal	Duplicate a principal by selecting the principal to duplicate in the Principal List, then clicking the Duplicate button.	"How to Duplicate a Principal" on page 148
Modify a principal	Modify a principal by selecting the principal to modify in the Principal List, then clicking the Modify button.  Note that you cannot modify a principal's name. To rename a principal, you must duplicate the principal, specify a new name for it, save it, and then delete the old principal.	"How to Modify a Principal" on page 148
Delete a principal	Delete a principal by selecting the principal to delete in the Principal List, then clicking the Delete button.	"How to Delete a Principal" on page 149
Set up defaults for creating new principals	Set up defaults for creating new principals by choosing Properties from the Edit menu.	"How to Set Up Defaults for Creating New Principals" on page 150
Modify the Kerberos administration privileges (kadm5.ac1 File)	<i>Command-line only.</i> The Kerberos administration privileges determine what operations a principal can perform on the Kerberos database, such as add and modify. You need to edit the <code>/etc/krb5/kadm5.ac1</code> file to modify the Kerberos administration privileges for each principal.	"How to Modify the Kerberos Administration Privileges" on page 151

## Automating the Creation of New Principals

Even though the SEAM Tool provides ease-of-use, it doesn't provide a way to automate the creation of new principals. Automation is especially useful if you need to add 10 or even 100 new principals in a short time. However, by using the `kadmin.local` command in a Bourne shell script, you can do just that.

The following shell script line is an example of how automate the creation of new principals:

```
sed -e 's/^\(.*\)$/ank +needchange -pw \1 \1/' < princnames |  
time /usr/sbin/kadmin.local> /dev/null
```

This example is split over two lines readability. The script reads in a file called `princnames` that contains principal names and their passwords, and adds them to the Kerberos database. You would have to create the `princnames` file, which contains a principal name and its password on each line, separated by one or more spaces. The `+needchange` option configures the principal so that the user is prompted for a new password during login with the principal for the first time. This practice helps to ensure that the passwords in the `princnames` file are not a security risk.

You can build more elaborate scripts. For example, your script could use the information in the name service to obtain the list of user names for the principal names. What you do and how you do it is determined by your site needs and your scripting expertise.

## ▼ How to View the List of Principals

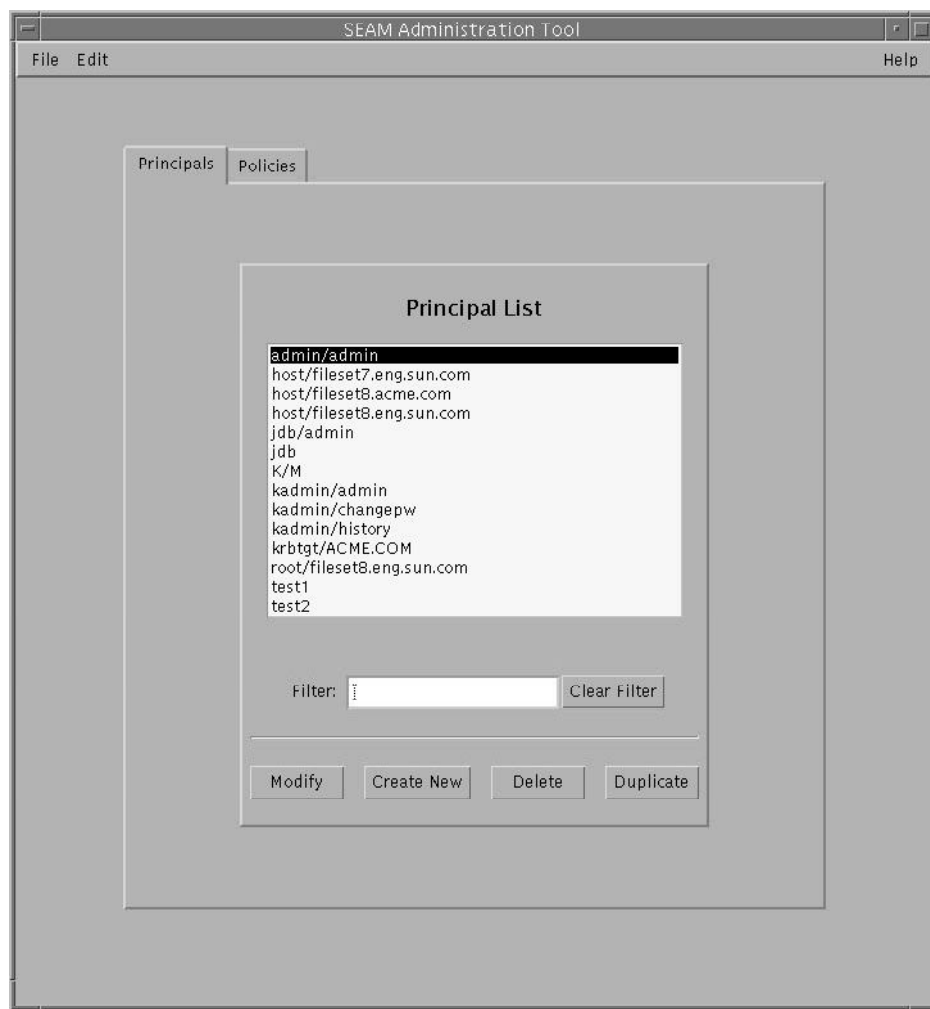
An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**

See “How to Start the SEAM Tool” on page 139 for details.

- 2. Click the Principals tab.**

The list of principals is displayed.



### 3. Display a specific principal or a sublist of principals.

Type a filter string in the Filter field, and press Return. If the filter succeeds, the list of principals that match the filter is displayed.

The filter string must consist of one or more characters. Because the filter mechanism is case sensitive, you need to use the appropriate uppercase and lowercase letters for the filter. For example, if you type the filter string *ge*, the filter mechanism displays only the principals with the *ge* string in them (for example, *george* or *edge*).

If you want to display the entire list of principals, click Clear Filter.

## Example—Viewing the List of Principals (Command Line)

In the following example, the `list_principals` command of `kadmin` is used to list all the principals that match `test*`. Wildcards can be used with the `list_principals` command.

```
kadmin: list_principals test*
test1@EXAMPLE.COM
test2@EXAMPLE.COM
kadmin: quit
```

## ▼ How to View a Principal's Attributes


An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**  
See "How to Start the SEAM Tool" on page 139 for details.
- 2. Click the Principals tab.**
- 3. Select the principal in the list that you want to view, then click Modify.**  
The Principal Basics panel that contains some of the principal's attributes is displayed.
- 4. Continue to click Next to view all the principal's attributes.**  
Three windows contain attribute information. Choose Context-Sensitive Help from the Help menu to get information about the various attributes in each window. Or, for all the principal attribute descriptions, go to "SEAM Tool Panel Descriptions" on page 162.
- 5. When you are finished viewing, click Cancel.**

## Example—Viewing a Principal's Attributes

The following example shows the first window when you are viewing the `jdb/admin` principal.





## Example—Viewing a Principal’s Attributes (Command Line)

In the following example, the `get_principal` command of `kadmin` is used to view the attributes of the `jdb/admin` principal.

```
kadmin: getprinc jdb/admin
Principal: jdb/admin@EXAMPLE.COM
Expiration date: Fri Aug 25 17:19:05 PDT 2000
Last password change: [never]
Password expiration date: Wed Apr 14 11:53:10 PDT 1999
Maximum ticket life: 1 day 16:00:00
```

```
Maximum renewable life: 1 day 16:00:00
Last modified: Thu Jan 14 11:54:09 PST 1999 (admin/admin@EXAMPLE.COM)
Last successful authentication: [never]
Last failed authentication: [never]
Failed password attempts: 0
Number of keys: 1
Key: vno 1, DES cbc mode with CRC-32, no salt
Attributes: REQUIRES_HW_AUTH
Policy: [none]
kadmin: quit
```

## ▼ How to Create a New Principal

An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**

See “How to Start the SEAM Tool” on page 139 for details.

---

**Note** – If you are creating a new principal that might need a new policy, you should create the new policy before you create the new principal. Go to “How to Create a New Policy” on page 158.

---

- 2. Click the Principals tab.**

- 3. Click New.**

The Principal Basics panel that contains some attributes for a principal is displayed.

- 4. Specify a principal name and a password.**

Both the principal name and password are mandatory.

- 5. Specify values for the principal’s attributes, and continue to click Next to specify more attributes.**

Three windows contain attribute information. Choose Context-Sensitive Help from the Help menu to get information about the various attributes in each window. Or, for all the principal attribute descriptions, go to “SEAM Tool Panel Descriptions” on page 162.


- 6. Click Save to save the principal, or click Done on the last panel.**

- 7. If needed, set up Kerberos administration privileges for the new principal in the /etc/krb5/kadm5.acl file.**

See “How to Modify the Kerberos Administration Privileges” on page 151 for more details.

## Example—Creating a New Principal

The following example shows the Principal Basics panel when a new principal called pak is created. The policy is set to testuser.



## Example—Creating a New Principal (Command Line)

In the following example, the `add_principal` command of `kadmin` is used to create a new principal called `pak`. The principal's policy is set to `testuser`.

```
kadmin: add_principal -policy testuser pak
Enter password for principal "pak@EXAMPLE.COM": <type the password>
```

```
Re-enter password for principal "pak@EXAMPLE.COM": <type the password again>
Principal "pak@EXAMPLE.COM" created.
kadmin: quit
```

## ▼ How to Duplicate a Principal

This procedure explains how to use all or some of the attributes of an existing principal to create a new principal. No command-line equivalent exists for this procedure.

- 1. If necessary, start the SEAM Tool.**

See “How to Start the SEAM Tool” on page 139 for details.

- 2. Click the Principals tab.**

- 3. Select the principal in the list that you want to duplicate, then click Duplicate.**

The Principal Basics panel is displayed. All the attributes of the selected principal are duplicated except for the Principal Name and Password fields, which are empty.

- 4. Specify a principal name and a password.**

Both the principal name and the password are mandatory. To make an exact duplicate of the principal you selected, click Save and skip to Step 7.

- 5. Specify different values for the principal’s attributes, and continue to click Next to specify more attributes.**

Three windows contain attribute information. Choose Context-Sensitive Help from the Help menu to get information about the various attributes in each window. Or, for all the principal attribute descriptions, go to “SEAM Tool Panel Descriptions” on page 162.

- 6. Click Save to save the principal, or click Done on the last panel.**

- 7. If needed, set up Kerberos administration privileges for the principal in `/etc/krb5/kadm5.acl` file.**

See “How to Modify the Kerberos Administration Privileges” on page 151 for more details.

## ▼ How to Modify a Principal

An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**

See “How to Start the SEAM Tool” on page 139 for details.

- 2. Click the Principals tab.**

**3. Select the principal in the list that you want to modify, then click Modify.**

The Principal Basics panel that contains some of the attributes for the principal is displayed.

**4. Modify the principal's attributes, and continue to click Next to modify more attributes.**

Three windows contain attribute information. Choose Context-Sensitive Help from the Help menu to get information about the various attributes in each window. Or, for all the principal attribute descriptions, go to "SEAM Tool Panel Descriptions" on page 162.

---

**Note** – You cannot modify a principal's name. To rename a principal, you must duplicate the principal, specify a new name for it, save it, and then delete the old principal.

---

**5. Click Save to save the principal, or click Done on the last panel.**

**6. Modify the Kerberos administration privileges for the principal in the /etc/krb5/kadm5.acl file.**

See "How to Modify the Kerberos Administration Privileges" on page 151 for more details.

## Example—Modifying a Principal's Password (Command Line)

In the following example, the `change_password` command of `kadmin` is used to modify the password for the `jdb` principal. The `change_password` command does not let you change the password to a password that is in the principal's password history.

```
kadmin: change_password jdb
Enter password for principal "jdb": <type the new password>
Re-enter password for principal "jdb": <type the password again>
Password for "jdb@EXAMPLE.COM" changed.
kadmin: quit
```

To modify other attributes for a principal, you must use the `modify_principal` command of `kadmin`.

## ▼ How to Delete a Principal

An example of the command-line equivalent follows this procedure.

**1. If necessary, start the SEAM Tool.**

See “How to Start the SEAM Tool” on page 139 for details.

**2. Click the Principals tab.**

**3. Select the principal in the list that you want to delete, then click Delete.**

After you confirm the deletion, the principal is deleted.

**4. Remove the principal from the Kerberos access control list (ACL) file,  
/etc/krb5/kadm5.acl.**

See “How to Modify the Kerberos Administration Privileges” on page 151 for more details.

## Example—Deleting a Principal (Command Line)

In the following example, the `delete_principal` command of `kadmin` is used to delete the `jdb` principal.

```
kadmin: delete_principal pak
Are you sure you want to delete the principal "pak@EXAMPLE.COM"? (yes/no): yes
Principal "pak@EXAMPLE.COM" deleted.
Make sure that you have removed this principal from all ACLs before reusing.
kadmin: quit
```

## ▼ How to Set Up Defaults for Creating New Principals


No command-line equivalent exists for this procedure.

**1. If necessary, start the SEAM Tool.**

See “How to Start the SEAM Tool” on page 139 for details.

**2. Choose Properties from the Edit Menu.**

The Properties window is displayed.



**3. Select the defaults that you want when you create new principals.**

Choose Context-Sensitive Help from the Help menu for information about the various attributes in each window.

**4. Click Save.**

## ▼ How to Modify the Kerberos Administration Privileges

Even though your site probably has many user principals, you usually want only a few users to be able to administer the Kerberos database. Privileges to administer the Kerberos database are determined by the Kerberos access control list (ACL) file, `kadm5.ac1`. The `kadm5.ac1` file enables you to allow or disallow privileges for individual principals. Or, you can use the `'*` wildcard in the principal name to specify privileges for groups of principals.

**1. Become superuser on the master KDC.**

## 2. Edit the `/etc/krb5/kadm5.ac1` file.

An entry in the `kadm5.ac1` file must have the following format:

```
principal privileges [principal-target]
```

*principal*

Specifies the principal to which the privileges are granted. Any part of the principal name can include the `'*` wildcard, which is useful for providing the same privileges for a group of principals. For example, if you want to specify all principals with the `admin` instance, you would use `*/admin@realm`.

Note that a common use of an `admin` instance is to grant separate privileges (such as administration access to the Kerberos database) to a separate Kerberos principal. For example, the user `jdb` might have a principal for his administrative use, called `jdb/admin`. This way, the user `jdb` obtains `jdb/admin` tickets only when he or she actually needs to use those privileges.

*privileges*

Specifies which operations can or cannot be performed by the principal. This field consists of a string of one or more of the following list of characters or their uppercase counterparts. If the character is uppercase (or not specified), then the operation is disallowed. If the character is lowercase, then the operation is permitted.

- a [Dis]allows the addition of principals or policies.
- d [Dis]allows the deletion of principals or policies.
- m [Dis]allows the modification of principals or policies.
- c [Dis]allows the changing of passwords for principals.
- i [Dis]allows inquiries to the Kerberos database.
- l [Dis]allows the listing of principals or policies in the Kerberos database.
- x or \* Allows all privileges (`admcil`).

*principal-target*

When a principal is specified in this field, the *privileges* apply to *principal* only when the *principal* operates on the *principal\_target*. Any part of the principal name can include the `'*` wildcard, which is useful to group principals.



## Example—Modifying the Kerberos Administration Privileges

The following entry in the `kadm5.ac1` file gives any principal in the `EXAMPLE.COM` realm with the `admin` instance all the privileges on the Kerberos database.

```
*/admin@EXAMPLE.COM *
```

The following entry in the `kadm5.ac1` file gives the `jdb@EXAMPLE.COM` principal the privilege to add, list, and inquire about any principal that has the `root` instance.

```
jdb@EXAMPLE.COM ali */root@EXAMPLE.COM
```

---

## Administering Policies

This section provides step-by-step instructions to administer policies with the SEAM Tool. This section also provides examples of equivalent command lines, when available.

### Administering Policies (Task Map)

Task	Description	For Instructions
View the list of policies	View the list of policies by clicking the Policies tab.	“How to View the List of Policies” on page 154
View a policy’s attributes	View a policy’s attributes by selecting the policy in the Policy List, then clicking the Modify button.	“How to View a Policy’s Attributes” on page 156
Create a new policy	Create a new policy by clicking the Create New button in the Policy List panel.	“How to Create a New Policy” on page 158
Duplicate a policy	Duplicate a policy by selecting the policy to duplicate in the Policy List, then clicking the Duplicate button.	“How to Duplicate a Policy” on page 160

Task	Description	For Instructions
Modify a policy	<p>Modify a policy by selecting the policy to modify in the Policy List, then clicking the Modify button.</p> <p>Note that you cannot modify a policy's name. To rename a policy, you must duplicate the policy, specify a new name for it, save it, and then delete the old policy.</p>	"How to Modify a Policy" on page 160
Delete a policy	Delete a policy by selecting the policy to delete in the Policy List, then clicking the Delete button.	"How to Delete a Policy" on page 161

## ▼ How to View the List of Policies

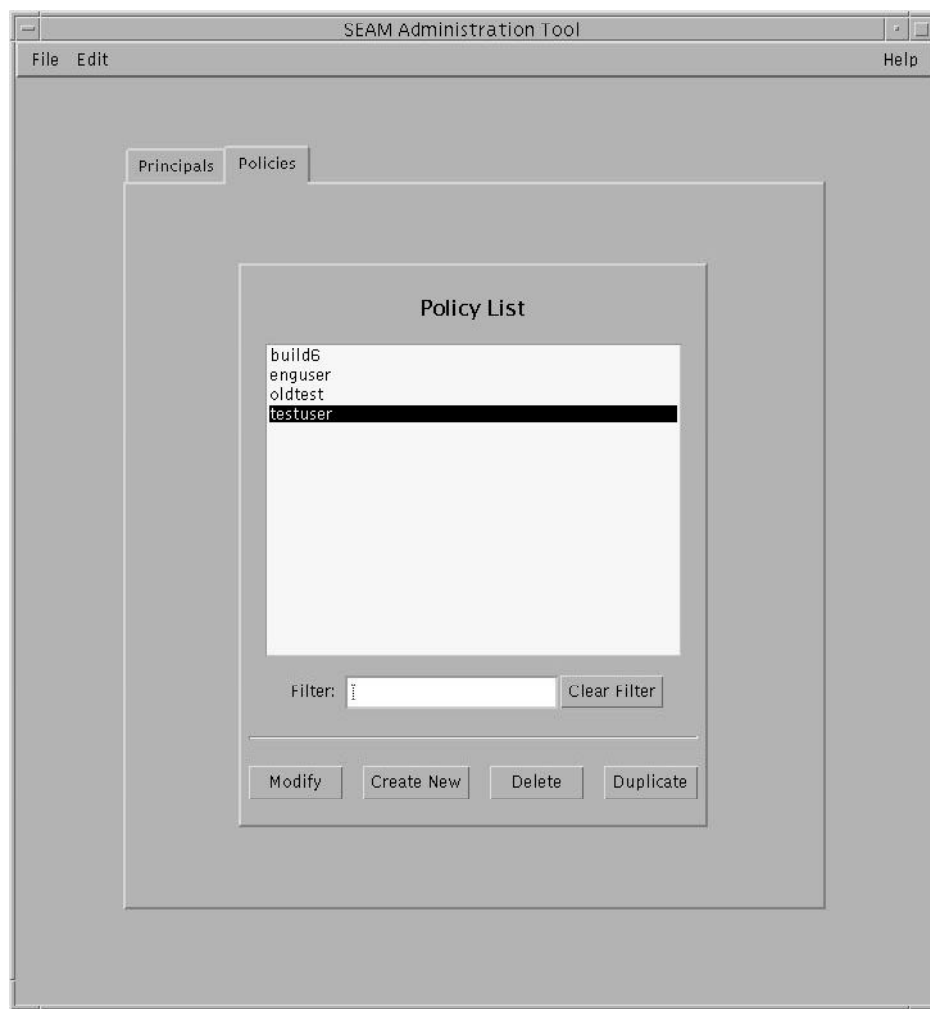
An example of the command-line equivalent follows this procedure.

**1. If necessary, start the SEAM Tool.**

See "How to Start the SEAM Tool" on page 139 for details.

**2. Click the Policies tab.**

The list of policies is displayed.



### 3. Display a specific policy or a sublist of policies.

Type a filter string in the Filter field, and press Return. If the filter succeeds, the list of policies that match the filter is displayed.

The filter string must consist of one or more characters. Because the filter mechanism is case sensitive, you need to use the appropriate uppercase and lowercase letters for the filter. For example, if you type the filter string `ge`, the filter mechanism displays only the policies with the `ge` string in them (for example, `george` or `edge`).

If you want to display the entire list of policies, click Clear Filter.

## Example—Viewing the List of Policies (Command Line)

In the following example, the `list_policies` command of `kadmin` is used to list all the policies that match `*user*`. Wildcards can be used with the `list_policies` command.

```
kadmin: list_policies *user*
testuser
enguser
kadmin: quit
```

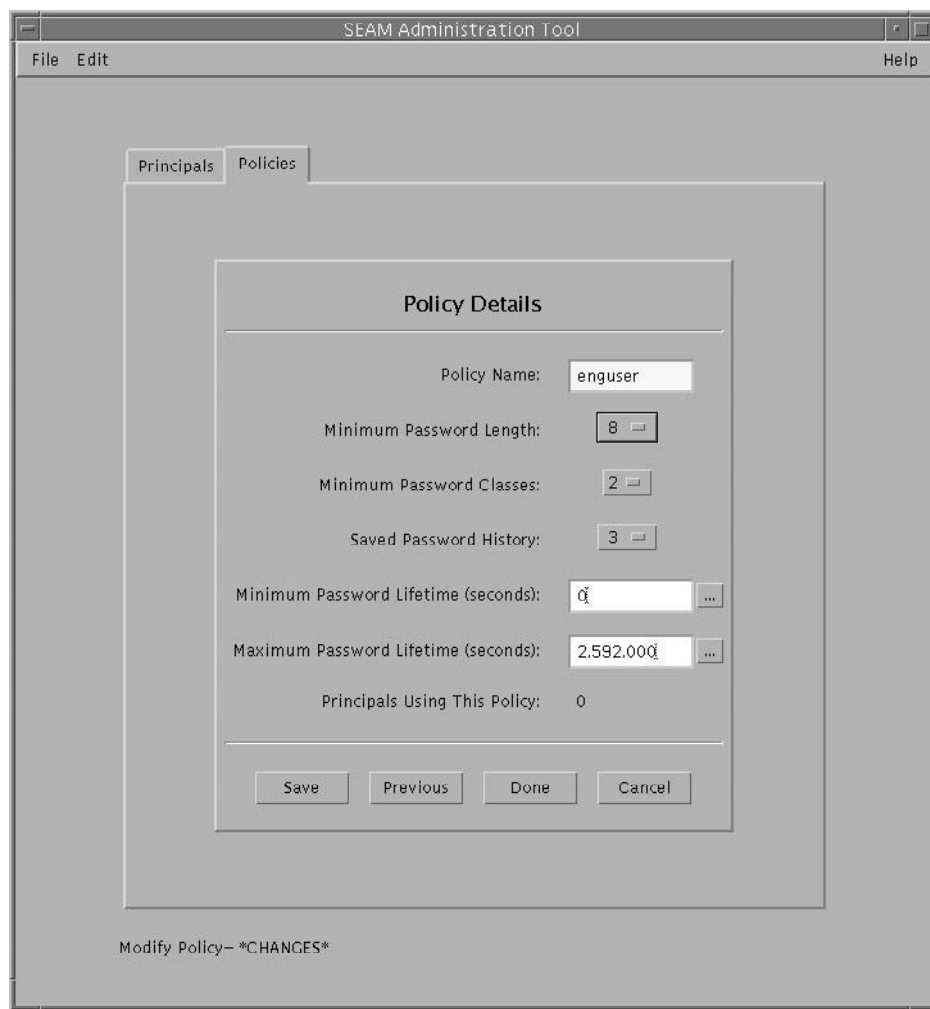
## ▼ How to View a Policy's Attributes

An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**  
See "How to Start the SEAM Tool" on page 139 for details.
- 2. Click the Policies tab.**
- 3. Select the policy in the list that you want to view, then click Modify.**  
The Policy Details panel is displayed.
- 4. When you are finished viewing, click Cancel.**

## Example—Viewing a Policy's Attributes

The following example shows the Policy Details panel when you are viewing the `test` policy.



## Example—Viewing a Policy’s Attributes (Command Line)

In the following example, the `get_policy` command of `kadmin` is used to view the attributes of the `enguser` policy.

```
kadmin: get_policy enguser
Policy: enguser
Maximum password life: 2592000
Minimum password life: 0
Minimum password length: 8
Minimum number of password character classes: 2
Number of old keys kept: 3
Reference count: 0
```

kadmin: **quit**

The reference count is the number of principals that use this policy.

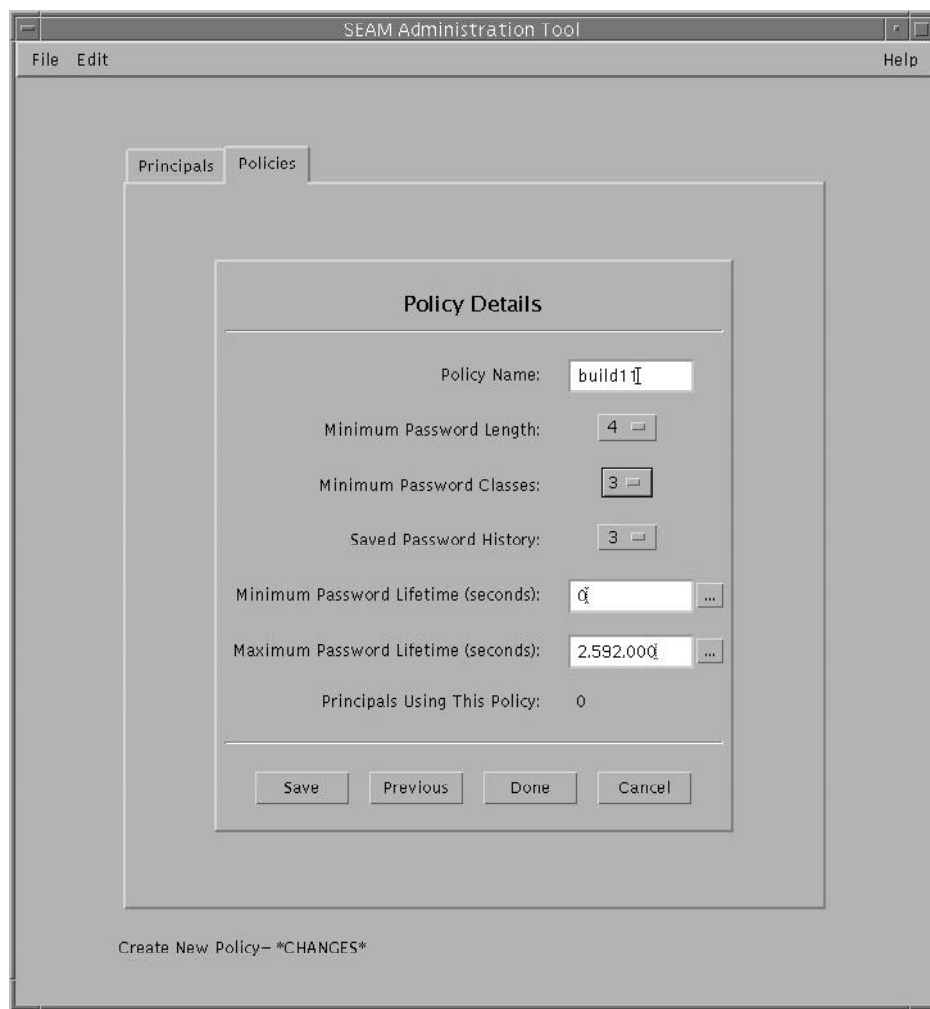
## ▼ How to Create a New Policy

An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**  
See “How to Start the SEAM Tool” on page 139 for details.
- 2. Click the Policies tab.**
- 3. Click New.**  
The Policy Details panel is displayed.
- 4. Specify a name for the policy in the Policy Name field.**  
The policy name is mandatory.
- 5. Specify values for the policy’s attributes.**  
Choose Context-Sensitive Help from the Help menu for information about the various attributes in this window. Or, go to Table 10–5 for all the policy attribute descriptions.
- 6. Click Save to save the policy, or click Done.**

### Example—Creating a New Policy

In the following example, a new policy called `build11` is created. The Minimum Password Classes is set to 3.



## Example—Creating a New Policy (Command Line)

In the following example, the `add_policy` command of `kadmin` is used to create the `build11` policy. This policy requires at least 3 character classes in a password.

```
$ kadmin
kadmin: add_policy -minclasses 3 build11
kadmin: quit
```

## ▼ How to Duplicate a Policy

This procedure explains how to use all or some of the attributes of an existing policy to create a new policy. No command-line equivalent exists for this procedure.

- 1. If necessary, start the SEAM Tool.**  
See “How to Start the SEAM Tool” on page 139 for details.
- 2. Click the Policies tab.**
- 3. Select the policy in the list that you want to duplicate, then click Duplicate.**  
The Policy Details panel is displayed. All the attributes of the selected policy are duplicated, except for the Policy Name field, which is empty.
- 4. Specify a name for the duplicated policy in the Policy Name field.**  
The policy name is mandatory. To make an exact duplicate of the policy you selected, click Save and skip to Step 6.
- 5. Specify different values for the policy’s attributes.**  
Choose Context-Sensitive Help from the Help menu for information about the various attributes in this window. Or, go to Table 10–5 for all the policy attribute descriptions.
- 6. Click Save to save the policy, or click Done.**

## ▼ How to Modify a Policy

An example of the command-line equivalent follows this procedure.

- 1. If necessary, start the SEAM Tool.**  
See “How to Start the SEAM Tool” on page 139 for details.
- 2. Click the Policies tab.**
- 3. Select the policy in the list that you want to modify, then click Modify.**  
The Policy Details panel is displayed.
- 4. Modify the policy’s attributes.**  
Choose Context-Sensitive Help from the Help menu for information about the various attributes in this window. Or, go to Table 10–5 for all the policy attribute descriptions.

---

**Note** – You cannot modify a policy’s name. To rename a policy, you must duplicate the policy, specify a new name for it, save it, and then delete the old policy.

---

- 5. Click Save to save the policy, or click Done.**



## Example—Modifying a Policy (Command Line)

In the following example, the `modify_policy` command of `kadmin` is used to modify the minimum length of a password to five characters for the `build11` policy.

```
$ kadmin
kadmin: modify_policy -minlength 5 build11
kadmin: quit
```

## ▼ How to Delete a Policy

An example of the command-line equivalent follows this procedure.

1. **If necessary, start the SEAM Tool.**  
See “How to Start the SEAM Tool” on page 139 for details.
2. **Click the Policies tab.**

---

**Note** – Before you delete a policy, you must cancel the policy from all principals that are currently using it. To do so, you need to modify the principals’ Policy attribute. The policy cannot be deleted if any principal is using it.

---

3. **Select the policy in the list that you want to delete, then click Delete.**  
After you confirm the deletion, the policy is deleted.

## Example—Deleting a Policy (Command Line)

In the following example, the `delete_policy` command of the `kadmin` command is used to delete the `build11` policy.

```
kadmin: delete_policy build11
Are you sure you want to delete the policy "build11"? (yes/no): yes
kadmin: quit
```

Before you delete a policy, you must cancel the policy from all principals that are currently using it. To do so, you need to use the `modify_principal -policy` command of `kadmin` on the affected principals. The `delete_policy` command fails if the policy is in use by a principal.

---

# SEAM Tool Reference

This section provides reference information for the SEAM Tool.

## SEAM Tool Panel Descriptions

This section provides descriptions for each principal and policy attribute that you can either specify or view in the SEAM Tool. The attributes are organized by the panel in which they are displayed.

**TABLE 10-2** Attributes for the Principal Basics Panel

Attribute	Description
Principal Name	The name of the principal (the <i>primary/instance</i> part of a fully-qualified principal name). A principal is a unique identity to which the KDC can assign tickets. If you are modifying a principal, you cannot edit a principal's name.
Password	The password for the principal. You can use the Generate Random Password button to create a random password for the principal.
Policy	A menu of available policies for the principal.
Account Expires	The date and time on which the principal's account expires. When the account expires, the principal can no longer get a ticket-granting ticket (TGT) and might be unable to log in.
Last Principal Change	The date on which information for the principal was last modified. (Read-only)
Last Changed By	The name of the principal that last modified the account for this principal. (Read-only)
Comments	Comments that are related to the principal (for example, "Temporary Account").

**TABLE 10-3** Attributes for the Principal Details Panel

Attribute	Description
Last Success	The date and time when the principal last logged in successfully. (Read-only)
Last Failure	The date and time when the last login failure for the principal occurred. (Read-only)
Failure Count	The number of times that there has been a login failure for the principal. (Read-only)
Last Password Change	The date and time when the principal's password was last changed. (Read-only)

**TABLE 10-3** Attributes for the Principal Details Panel (Continued)

Attribute	Description
Password Expires	The date and time when the principal's current password expires.
Key Version	The key version number for the principal. This attribute is normally changed only when a password has been compromised.
Maximum Lifetime (seconds)	The maximum length of time for which a ticket can be granted for the principal (without renewal).
Maximum Renewal (seconds)	The maximum length of time for which an existing ticket can be renewed for the principal.

**TABLE 10-4** Attributes of the Principal Flags Panel

Attribute (Radio Buttons)	Description
Disable Account	When checked, prevents the principal from logging in. This attribute provides an easy way to temporarily freeze a principal account.
Require Password Change	When checked, expires the principal's current password, which forces the user to use the <code>kpasswd</code> command to create a new password. This attribute is useful if there is a security breach, and you need to make sure that old passwords are replaced.
Allow Postdated Tickets	When checked, allows the principal to obtain postdated tickets.  For example, you might need to use postdated tickets for <code>cron</code> jobs that must run after hours, but you cannot obtain tickets in advance because of short ticket lifetimes.
Allow Forwardable Tickets	When checked, allows the principal to obtain forwardable tickets.  Forwardable tickets are tickets that are forwarded to the remote host to provide a single-sign-on session. For example, if you are using forwardable tickets and you authenticate yourself through <code>ftp</code> or <code>rsh</code> , then other services, such as NFS services, are available without your being prompted for another password.
Allow Renewable Tickets	When checked, allows the principal to obtain renewable tickets.  A principal can automatically extend the expiration date or time of a ticket that is renewable (rather than having to get a new ticket after the first ticket expires). Currently, the NFS service is the ticket service that can renew tickets.
Allow Proxiable Tickets	When checked, allows the principal to obtain proxiable tickets.  A proxiable ticket is a ticket that can be used by a service on behalf of a client to perform an operation for the client. With a proxiable ticket, a service can take on the identity of a client and obtain a ticket for another service, but the service cannot obtain a ticket-granting ticket.

**TABLE 10-4** Attributes of the Principal Flags Panel (Continued)

Attribute (Radio Buttons)	Description
Allow Service Tickets	<p>When checked, allows service tickets to be issued for the principal.</p> <p>You should not allow service tickets to be issued for the <code>kaadmin/hostname</code> and <code>changeppw/hostname</code> principals. This practice ensures that these principals can only update the KDC database.</p>
Allow TGT-Based Authentication	<p>When checked, allows the service principal to provide services to another principal. More specifically, this attribute allows the KDC to issue a service ticket for the service principal.</p> <p>This attribute is valid only for service principals. When unchecked, service tickets cannot be issued for the service principal.</p>
Allow Duplicate Authentication	<p>When checked, allows the user principal to obtain service tickets for other user principals.</p> <p>This attribute is valid only for user principals. When unchecked, the user principal can still obtain service tickets for service principals, but not for other user principals.</p>
Required Preauthentication	<p>When checked, the KDC will not send a requested ticket-granting ticket (TGT) to the principal until the KDC can authenticate (through software) that the principal is really the principal that is requesting the TGT. This preauthentication is usually done through an extra password, for example, from a DES card.</p> <p>When unchecked, the KDC does not need to preauthenticate the principal before the KDC sends a requested TGT to the principal.</p>
Required Hardware Authentication	<p>When checked, the KDC will not send a requested ticket-granting ticket (TGT) to the principal until the KDC can authenticate (through hardware) that it is really the principal that is requesting the TGT. Hardware preauthentication can occur, for example, on a Java ring reader.</p> <p>When unchecked, the KDC does not need to preauthenticate the principal before the KDC sends a requested TGT to the principal.</p>

**TABLE 10-5** Attributes for the Policy Basics Pane

Attribute	Description
Policy Name	<p>The name of the policy. A policy is a set of rules that govern a principal's password and tickets.</p> <p>If you are modifying a policy, you cannot edit a policy's name.</p>
Minimum Password Length	<p>The minimum length for the principal's password.</p>

**TABLE 10-5** Attributes for the Policy Basics Pane (Continued)

Attribute	Description
Minimum Password Classes	The minimum number of different character types that are required in the principal's password.  For example, a minimum classes value of 2 means that the password must have at least two different character types, such as letters and numbers (hi2mom). A value of 3 means that the password must have at least three different character types, such as letters, numbers, and punctuation (hi2mom!). And so on.  A value of 1 sets no restriction on the number of password character types.
Saved Password History	The number of previous passwords that have been used by the principal, and a list of the previous passwords that cannot be reused.
Minimum Password Lifetime (seconds)	The minimum time that the password must be used before it can be changed.
Maximum Password Lifetime (seconds)	The maximum time that the password can be used before it must be changed.
Principals Using This Policy	The number of principals to which this policy currently applies. (Read-only)

## Using the SEAM Tool With Limited Kerberos Administration Privileges

All features of the SEAM Administration Tool are available if your admin principal has all the privileges to administer the Kerberos database. But it is possible to have limited privileges, such as being allowed to view only the list of principals or to change a principal's password. With limited Kerberos administration privileges, you can still use the SEAM Tool. However, various parts of the SEAM Tool will change based on the Kerberos administration privileges that you do not have. Table 10-6 shows how the SEAM Tool changes based on your Kerberos administration privileges.

The most visual change to the SEAM Tool occurs when you don't have the list privilege. Without the list privilege, the List panels do not display the list of principals and policies for you to manipulate. Instead, you must use the Name field in the List panels to specify a principal or policy that you want to manipulate.

If you login to the SEAM Tool, and you do not have sufficient privileges to perform tasks with it, the following message displays and you are sent back to the SEAM Administration Login window:

```
Insufficient privileges to use gkadmin: ADMCIL. Please try using another principal.
```

To change the privileges for a principal to administer the Kerberos database, go to "How to Modify the Kerberos Administration Privileges" on page 151.

**TABLE 10-6** Using SEAM Tool With Limited Kerberos Administration Privileges

Disallowed Privilege	Change the SEAM Tool
a (add)	The Create New and Duplicate buttons are unavailable in the Principal List and Policy List panels. Without the add privilege, you cannot create new principals or policies or duplicate them.
d (delete)	The Delete button is unavailable in the Principal List and Policy List panels. Without the delete privilege, you cannot delete principal or policies.
m (modify)	The Modify button is unavailable in the Principal List and Policy List panels. Without the modify privilege, you cannot modify principal or policies.  Also, with the Modify button unavailable, you cannot modify a principal's password, even if you have the change password privilege.
c (change password)	The Password field in the Principal Basics panel is read-only and cannot be changed. Without the change password privilege, you cannot modify a principal's password.  Note that even if you have the change password privilege, you must also have the modify privilege to change a principal's password.
i (inquiry to database)	The Modify and Duplicate buttons are unavailable in the Principal List and Policy List panels. Without the inquiry privilege, you cannot modify or duplicate a principal or policy.  Also, with the Modify button unavailable, you cannot modify a principal's password, even if you have the change password privilege.
l (list)	The list of principals and policies in the List panels are unavailable. Without the list privilege, you must use the Name field in the List panels to specify the principal or policy that you want to manipulate.

---

## Administering Keytab Files

Every host that provides a service must have a local file, called a *keytab* (short for key table). The keytab contains the principal for the appropriate service, called a *service key*. A service key is used by a service to authenticate itself to the KDC and is known only by Kerberos and the service itself. For example, if you have a Kerberized NFS server, that server must have a keytab file that contains its `nfs` service principal.

To add a service key to a keytab file, you add the appropriate service principal to a host's keytab file by using the `ktadd` command of `kadmin`. Because you are adding a service principal to a keytab file, the principal must already exist in the Kerberos database so that `kadmin` can verify its existence. On the master KDC, the keytab file is located at `/etc/krb5/kadm5.keytab`, by default. On application servers that provide Kerberized services, the keytab file is located at `/etc/krb5/krb5.keytab`, by default.

A keytab is analogous to a user's password. Just as it is important for users to protect their passwords, it is equally important for application servers to protect their keytab files. You should always store keytab files on a local disk, and make them readable only by the `root` user. Also, you should never send a keytab file over an unsecured network.

There is also a special instance to add a `root` principal to a host's keytab file. If you want a user on the SEAM client to mount Kerberized NFS file systems that use Kerberos authentication automatically, you must add the client's `root` principal to the client's keytab file. Otherwise, users must use the `kinit` command as `root` to obtain credentials for the client's `root` principal whenever they want to mount a Kerberized NFS file system, even when they are using the automounter. See "Setting Up Root Authentication to Mount NFS File Systems" on page 109 for detailed information.

---

**Note** – When you set up a master KDC, you need to add the `kadmin` and `changepw` principals to the `kadm5.keytab` file. This step enables the KDC to decrypt administrators' Kerberos tickets to determine whether it should give the administrators access to the database.

---

Another command that you can use to administer keytab files is the `ktutil` command. `ktutil` is an interactive command that enables you to manage a local host's keytab file without having Kerberos administration privileges, because `ktutil` doesn't interact with the Kerberos database as `kadmin` does. So, after a principal is added to a keytab file, you can use `ktutil` to view the keylist in a keytab file or to temporarily disable authentication for a service.

## Administering Keytabs Task Map

Task	Description	For Instructions
Add a service principal to a keytab file	Use the <code>ktadd</code> command of <code>kadmin</code> to add a service principal to a keytab file.	"How to Add a Service Principal to a Keytab File" on page 168

Task	Description	For Instructions
Remove a service principal from a keytab file	Use the <code>ktremove</code> command of <code>kadmin</code> to remove a service from a keytab file.	"How to Remove a Service Principal From a Keytab File" on page 169
Display the keylist (Principals) in a keytab file	Use the <code>ktutil</code> command to display the keylist in a keytab file.	"How to Display the Keylist (Principals) in a Keytab File" on page 170
Temporarily disable authentication for a service on a host	This procedure is a quick way to temporarily disable authentication for a service on a host without having to have <code>kadmin</code> privileges.  Before you use <code>ktutil</code> to delete the service principal from the server's keytab file, copy the original keytab file to a temporary location. When you want to enable the service again, copy the original keytab file back to its proper location.	"How to Temporarily Disable Authentication for a Service on a Host" on page 171

## ▼ How to Add a Service Principal to a Keytab File

1. **Make sure that the principal already exists in the Kerberos database.**

See "How to View the List of Principals" on page 142 for more information.

2. **Become superuser on the host that needs a principal added to its keytab file.**

3. **Start the `kadmin` command.**

```
# /usr/sbin/kadmin
```

4. **Add a principal to a keytab file by using the `ktadd` command.**

```
kadmin: ktadd [-k keytab] [-q] [principal | -glob principal-exp]
```

<code>-k <i>keytab</i></code>	Specifies the keytab file. By default, <code>/etc/krb5/krb5.keytab</code> is used.
<code>-q</code>	Displays less verbose information.
<code><i>principal</i></code>	Specifies the principal to be added to the keytab file. You can add the following service principals: <code>host</code> , <code>root</code> , <code>nfs</code> , and <code>ftp</code> .
<code>-glob <i>principal-exp</i></code>	Specifies the principal expressions. All principals that match the principal are added to the keytab file. The rules for principal expression are the same as for the <code>list_principals</code> command of <code>kadmin</code> .

5. **Quit the `kadmin` command.**



```
kadmin: quit
```

## Example—Adding a Service Principal to a Keytab File

In the following example, the `kadmin/admin` and `kadmin/changepw` principals are added to a master KDC's keytab file. For this example, the keytab file must be the file that is specified in the `kdc.conf` file.

```
kdc1 # /usr/sbin/kadmin.local
kadmin.local: ktadd -k /etc/krb5/kadm5.keytab kadmin/admin kadmin/changepw
Entry for principal kadmin/admin@EXAMPLE.COM with kvno 3, encryption type DES-CBC-CRC
added to keytab WRFILE:/etc/krb5/kadm5.keytab.
Entry for principal kadmin/changepw@EXAMPLE.COM with kvno 3, encryption type DES-CBC-CRC
added to keytab WRFILE:/etc/krb5/kadm5.keytab.
kadmin.local: quit
```

In the following example, `denver`'s host principal is added to `denver`'s keytab file, so that the KDC can authenticate `denver`'s network services.

```
denver # /usr/sbin/kadmin
kadmin: ktadd host/denver@example.com@EXAMPLE.COM
kadmin: Entry for principal host/denver@example.com@EXAMPLE.COM with kvno 2,
encryption type DES-CBC-CRC added to keytab WRFILE:/etc/krb5/krb5.keytab.
kadmin: quit
```

## ▼ How to Remove a Service Principal From a Keytab File

1. Become superuser on the host with a service principal that must be removed from its keytab file.
2. Start the `kadmin` command.

```
# /usr/sbin/kadmin
```

3. (Optional) To display the current list of principals (keys) in the keytab file, use the `ktutil` command.

See "How to Display the Keylist (Principals) in a Keytab File" on page 170 for detailed instructions.

4. Remove a principal from the keytab file by using the `ktremove` command.

```
kadmin: ktrremove [-k keytab] [-q] principal [kvno | all | old ]
```

`-k keytab`

Specifies the keytab file. By default, `/etc/krb5/krb5.keytab` is used.

<code>-q</code>	Displays less verbose information.
<i>principal</i>	Specifies the principal to be removed from the keytab file.
<i>kvno</i>	Removes all entries for the specified principal whose key version number matches <i>kvno</i> .
<code>all</code>	Removes all entries for the specified principal.
<code>old</code>	Removes all entries for the specified principal, except those principals with the highest key version number.

#### 5. Quit the `kadmin` command.

```
kadmin: quit
```

### Example—Removing a Service Principal From a Keytab

In the following example, `denver`'s host principal is removed from `denver`'s keytab file.

```
denver # /usr/sbin/kadmin
kadmin: ktremove host/denver.example.com@EXAMPLE.COM
kadmin: Entry for principal host/denver.example.com@EXAMPLE.COM with kvno 3
        removed from keytab WRFILE:/etc/krb5/krb5.keytab.
kadmin: quit
```

## ▼ How to Display the Keylist (Principals) in a Keytab File

### 1. Become superuser on the host with the keytab file.

---

**Note** – Although you can create keytab files that are owned by other users, the default location for the keytab file requires `root` ownership.

---

### 2. Start the `ktutil` command.

```
# /usr/bin/ktutil
```

### 3. Read the keytab file into the keylist buffer by using the `read_kt` command.

```
ktutil: read_kt keytab
```

### 4. Display the keylist buffer by using the `list` command.

```
ktutil: list
```

The current keylist buffer is displayed.

5. **Quit the `ktutil` command.**

```
ktutil: quit
```

## Example—Displaying the Keylist (Principals) in a Keytab File

The following example displays the keylist in the `/etc/krb5/krb5.keytab` file on the `denver` host.

```
denver # /usr/bin/ktutil
ktutil: read_kt /etc/krb5/krb5.keytab
ktutil: list
slot KVNO Principal
-----
1    5 host/denver@EXAMPLE.COM
ktutil: quit
```

## ▼ How to Temporarily Disable Authentication for a Service on a Host

At times, you might need to temporarily disable the authentication mechanism for a service, such as `rlogin` or `ftp`, on a network application server. For example, you might want to stop users from logging in to a system while you are performing maintenance procedures. The `ktutil` command enables you to accomplish this task by removing the service principal from the server's keytab file, without requiring `kadmin` privileges. To enable authentication again, you just need to copy the original keytab file that you saved back to its original location.

---

**Note** – By default, most services are set up to require authentication. If a service is not set up to require authentication, then the service will still work, even if you disable authentication for the service.

---

1. **Become superuser on the host with the keytab file.**

---

**Note** – Although you can create keytab files that are owned by other users, the default location for the keytab file requires `root` ownership.

---

2. **Save the current keytab file to a temporary file.**

3. **Start the `ktutil` command.**

```
# /usr/bin/ktutil
```

4. Read the keytab file into the keylist buffer by using the `read_kt` command.

```
ktutil: read_kt keytab
```

5. Display the keylist buffer by using the `list` command.

```
ktutil: list
```

The current keylist buffer is displayed. Note the slot number for the service that you want to disable.

6. To temporarily disable a host's service, remove the specific service principal from the keylist buffer by using the `delete_entry` command.

```
ktutil: delete_entry slot-number
```

In this example, *slot-number* specifies the slot number of the service principal to be deleted, which is displayed by the `list` command.

7. Write the keylist buffer to the keytab file by using the `write_kt` command.

```
ktutil: write_kt keytab
```

8. Quit the `ktutil` command.

```
ktutil: quit
```

9. When you want to re-enable the service, copy the temporary (original) keytab file back to its original location.

## Example—Temporarily Disabling a Service on a Host

In the following example, the host service on the `denver` host is temporarily disabled. To enable the host service back on `denver`, you would copy the `krb5.keytab.temp` file to the `/etc/krb5/krb5.keytab` file.

```
denver # cp /etc/krb5/krb5.keytab /etc/krb5/krb5.keytab.temp
denver # /usr/bin/ktutil
      ktutil:read_kt /etc/krb5/krb5.keytab
      ktutil:list
slot KVNO Principal
-----
  1    8 root/denver@EXAMPLE.COM
  2    5 host/denver@EXAMPLE.COM
      ktutil:delete_entry 2
      ktutil:list
slot KVNO Principal
-----
  1    8 root/denver@EXAMPLE.COM
      ktutil:write_kt /etc/krb5/krb5.keytab
      ktutil:quit
```

## Using SEAM (Tasks)

---

This chapter is intended for anyone on a system with SEAM installed on it. This chapter includes information on tickets: obtaining, viewing, and destroying them. This chapter also includes information on choosing or changing a Kerberos password.

This is a list of the information in this chapter:

- “Ticket Management” on page 173
- “Password Management” on page 176

For an overview of SEAM, see Chapter 6.

---

## Ticket Management

This section explains how to obtain, view, and destroy tickets. For an introduction to tickets, see “How SEAM Works” on page 72.

### Do You Need to Worry About Tickets?

With SEAM 1.0 or 1.0.1 installed, Kerberos is built into the `login` command, and you will obtain tickets automatically when you log in.

Most of the Kerberized commands also automatically destroy your tickets when they exit. However, you might want to explicitly destroy your Kerberos tickets with `kdestroy` when you are finished with them, just to be sure. See “How to Destroy Tickets” on page 176 for more information on `kdestroy`.

For information on ticket lifetimes, see “Ticket Lifetimes” on page 187.

## ▼ How to Create a Ticket

Normally, a ticket is created automatically when you log in, and you need not do anything special to obtain a ticket. However, you might need to create a ticket if your ticket expires.

To create a ticket, use the `kinit` command.

```
% /usr/bin/kinit
```

`kinit` prompts you for your password. For the full syntax of the `kinit` command, see the `kinit(1)` man page.

### Example—Creating a Ticket

This example shows a user, `jennifer`, creating a ticket on her own system:

```
% kinit
Password for jennifer@ENG.EXAMPLE.COM: <type password>
```

Here, the user  `david`  creates a ticket that is valid for three hours with the `-l` option:

```
% kinit -l 3h david@EXAMPLE.ORG
Password for david@EXAMPLE.ORG: <type password>
```

This example shows the user  `david`  creating a forwardable ticket (with the `-f` option) for himself. With this forwardable ticket, he can, for example, log in to a second system.

```
% kinit -f david@EXAMPLE.ORG
Password for david@EXAMPLE.ORG: <type password>
```

For more on how forwarding tickets works, see “Types of Tickets” on page 186.

## ▼ How to View Tickets

Not all tickets are alike. One ticket might be, for example, *forwardable*; another ticket might be *postdated*; while a third ticket might be both forwardable and postdated. You can see which tickets you have, and what their attributes are, by using the `klist` command with the `-f` option:

```
% /usr/bin/klist -f
```

The following symbols indicate the attributes that are associated with each ticket, as displayed by `klist`:

F	Forwardable
f	Forwarded
P	Proxiabile
p	Proxy
D	Postdateable
d	Postdated
R	Renewable
I	Initial
i	Invalid

“Types of Tickets” on page 186 describes the various attributes that a ticket can have.

## Example—Viewing Tickets

This example shows that the user `jennifer` has an *initial* ticket, which is *forwardable* (F) and *postdated* (d), but not yet validated (i):

```
% /usr/bin/klist -f
Ticket cache: /tmp/krb5cc_74287
Default principal: jennifer@ENG.EXAMPLE.COM

Valid starting          Expires                Service principal
09 Mar 99 15:09:51    09 Mar 99 21:09:51    nfs/EXAMPLE.SUN.COM@EXAMPLE.SUN.COM
        renew until 10 Mar 99 15:12:51, Flags: Fdi
```

The following example shows that the user `david` has two tickets that were *forwarded* (f) to his host from another host. The tickets are also *forwardable* (F):

```
% klist -f
Ticket cache: /tmp/krb5cc_74287
Default principal: david@EXAMPLE.SUN.COM

Valid starting          Expires                Service principal
07 Mar 99 06:09:51    09 Mar 99 23:33:51    host/EXAMPLE.COM@EXAMPLE.COM
        renew until 10 Mar 99 17:09:51, Flags: fF

Valid starting          Expires                Service principal
08 Mar 99 08:09:51    09 Mar 99 12:54:51    nfs/EXAMPLE.COM@EXAMPLE.COM
        renew until 10 Mar 99 15:22:51, Flags: fF
```

## ▼ How to Destroy Tickets

Usually, tickets are destroyed automatically when the commands that created them exit. However, you might want to explicitly destroy your Kerberos tickets when you are finished with them, just to be sure. Tickets can be stolen. If tickets are stolen, the person who has stolen them can use them until they expire (although stolen tickets must be decrypted).

To destroy your tickets, use the `kdestroy` command.

```
% /usr/bin/kdestroy
```

`kdestroy` destroys *all* your tickets. You cannot use this command to selectively destroy a particular ticket.

If you are going to be away from your system and are concerned about an intruder using your permissions, you should use either `kdestroy` or a screen saver that locks the screen.

---

**Note** – One way to help ensure that your tickets are always destroyed is to add the `kdestroy` command to the `.logout` file in your home directory.

In instances where the PAM module has been configured (which is the default and usual case), tickets are destroyed automatically upon logout. So, adding a call to `kdestroy` to your `.login` file is not necessary. However, if the PAM module has not been configured, or if you don't know whether it has been, you might want to add `kdestroy` to your `.login` file to ensure that your tickets are destroyed when you exit your system.

---

---

## Password Management

With SEAM installed, you now have two passwords: your regular Solaris password, and a Kerberos password. You can make both passwords the same, or they can be different.

Non-Kerberized commands, such as `login`, are typically set up through PAM to authenticate with both Kerberos and UNIX. If you have different passwords, you must provide both passwords to log on with the appropriate authentication. However, if both passwords are the same, the first password you enter for UNIX is also accepted by Kerberos.



Unfortunately, using the same password for both Kerberos and UNIX can compromise security. That is, if someone discovers your Kerberos password, then your UNIX password is no longer a secret. However, using the same passwords for UNIX and Kerberos is still more secure than in a site without Kerberos, because passwords in a Kerberos environment are not sent across the network. Usually, your site will have a policy to help you determine your options.

Your Kerberos password is the only way Kerberos can verify your identity. If someone discovers your Kerberos password, Kerberos security becomes meaningless, because that person can masquerade as you. That person can send email that comes from “you,” read, edit, or delete your files, or log into other hosts as you. No one will be able to tell the difference. For this reason, it is vital that you choose a good password and keep it secret. You should *never* reveal your password to anyone else, not even your system administrator. Additionally, you should change your password frequently, particularly any time that you believe someone might have discovered it.

## Advice on Choosing a Password

Your password can include almost any character that you can type. The main exceptions are the Control keys and the Return key. A good password is a password that you can remember readily, but which no one else can easily guess. Examples of bad passwords include the following:

- Words that can be found in a dictionary
- Any common or popular name
- The name of a famous person or character
- Your name or user name in any form (for example: your name spelled backward, repeated twice, and so forth)
- A spouse’s name, child’s name, or pet’s name
- Your birth date or a relative’s birth date
- Your social security number, driver’s license number, passport number, or other similar identifying number
- Any sample password that appears in this manual or any other manual

A good password is at least eight characters long. Moreover, a password should include a mix of characters, such as uppercase and lowercase letters, numbers, and punctuation marks. Examples of passwords that would be good if they didn’t appear in this manual include the following:

- Acronyms, such as “I2LMHInSF” (which is recalled as “I too left my heart in San Francisco”)
- Easy-to-pronounce nonsense words, such as “WumpaBun” or “WangDangdoodle!”
- Deliberately misspelled phrases, such as “6o’cluck” or “RrriotGrrrlsRrrule!”



---

**Caution** – Don't use these examples. Passwords that appear in manuals are the first passwords that an intruder will try.

---

## Changing Your Password

You can change your Kerberos password in two ways:

- With the usual UNIX `passwd` command. With SEAM installed, the Solaris `passwd` command also automatically prompts for a new Kerberos password.

The advantage of using `passwd` instead of `kpasswd` is that you can set both passwords (UNIX and Kerberos) at the same time. However, you generally do not *have* to change both passwords with `passwd`. Often, you can change only your UNIX password and leave the Kerberos password untouched, or vice-versa.

---

**Note** – The behavior of `passwd` depends on how the PAM module is configured. You might be required to change both passwords in some configurations. For some sites, the UNIX password must be changed, while other sites require the Kerberos password to change.

---

- With the `kpasswd` command. `kpasswd` is very similar to `passwd`. One difference is that `kpasswd` changes only Kerberos passwords. You must use `passwd` if you want to change your UNIX password.

Another difference is that `kpasswd` can change a password for a Kerberos principal that is not a valid UNIX user. For example, `david/admin` is a Kerberos principal, but not an actual UNIX user, so you must use `kpasswd` instead of `passwd`.

After you change your password, it takes some time for the change to propagate through a system (especially over a large network). Depending on how your system is set up, this delay might take anywhere from a few minutes to an hour or more. If you need to get new Kerberos tickets shortly after you change your password, try the new password first. If the new password doesn't work, try again using the old password.

Kerberos V5 allows system administrators to set criteria about allowable passwords for each user. Such criteria is defined by the *policy* set for each user (or by a default policy). See "Administering Policies" on page 153 for more on policies.

For example, suppose that user `jennifer`'s policy (call it `jenpol`) mandates that passwords be at least eight letters long and include a mix of at least two kinds of characters. `kpasswd` will therefore reject an attempt to use "sloth" as a password.

```
% kpasswd
kpasswd: Changing password for jennifer@ENG.EXAMPLE.COM.
```

```

Old password:          <jennifer types her existing password>
kpasswd: jennifer@ENG.EXAMPLE.COM's password is controlled by
the policy jenpol
which requires a minimum of 8 characters from at least 2 classes
(the five classes are lowercase, uppercase, numbers, punctuation,
and all other characters).
New password:         <jennifer types 'sloth'>
New password (again): <jennifer re-types 'sloth'>
kpasswd: New password is too short.
Please choose a password which is at least 4 characters long.

```

Here, jennifer uses “slothrop49” as a password. “slothrop49” meets the criteria, because it is over eight letters long and contains two different kinds of characters (numbers and lowercase letters).

```

% kpasswd
kpasswd: Changing password for jennifer@ENG.EXAMPLE.COM.
Old password:          <jennifer types her existing password>
kpasswd: jennifer@ENG.EXAMPLE.COM's password is controlled by
the policy jenpol
which requires a minimum of 8 characters from at least 2 classes
(the five classes are lowercase, uppercase, numbers, punctuation,
and all other characters).
New password:         <jennifer types 'slothrop49'>
New password (again): <jennifer re-types 'slothrop49'>
Kerberos password changed.

```

## Examples—Changing Your Password

In the following example, user david changes both his UNIX password and Kerberos password with passwd.

```

% passwd
passwd: Changing password for david
Enter login (NIS+) password:          <type the current UNIX password>
New password:                          <type the new UNIX password>
Re-enter password:                     <confirm the new UNIX password>
Old KRB5 password:                     <type the current Kerberos password>
New KRB5 password:                     <type the new Kerberos password>
Re-enter new KRB5 password:            <confirm the new Kerberos password>

```

In the preceding example passwd asks for both the UNIX password and the Kerberos password. However, if *try\_first\_pass* is set in the PAM module, the Kerberos password is automatically set to the UNIX password. This is the default configuration. In that case, user david must use kpasswd to set his Kerberos password to something else, as shown next.

This example shows user david changing only his Kerberos password with kpasswd.

```

% kpasswd
kpasswd: Changing password for david@ENG.EXAMPLE.COM.
Old password:          <type the current Kerberos password>

```

```
New password:                <type the new Kerberos password>
New password (again):        <confirm the new Kerberos password>
Kerberos password changed.
```

In this example, user david changes the password for the Kerberos principal david/admin (which is not a valid UNIX user). He must use kpasswd.

```
% kpasswd david/admin
kpasswd: Changing password for david/admin.
Old password:                <type the current Kerberos password>
New password:                <type the new Kerberos password>
New password (again):        <type the new Kerberos password>
Kerberos password changed.
```

---

## SEAM (Reference)

---

This chapter lists many of the files, commands, and daemons that are part of the SEAM product. In addition, this chapter provides detailed information about how the Kerberos authentication system works.

This is a list of the reference information in this chapter.

- “SEAM Files” on page 181
- “SEAM Commands” on page 183
- “SEAM Daemons” on page 184
- “SEAM Terminology” on page 184
- “How the Authentication System Works” on page 189
- “Gaining Access to a Service Using SEAM” on page 190
- “Using the `gsscred` Table” on page 193

---

## SEAM Files

**TABLE 12-1** SEAM Files

File Name	Description
<code>~/ .gkadmin</code>	Default values for creating new principals in the SEAM Administration Tool
<code>~/ .k5login</code>	List of principals to grant access to a Kerberos account
<code>/etc/init.d/kdc</code>	init script to start or stop <code>krb5kdc</code>
<code>/etc/init.d/kdc.master</code>	init script to start or stop <code>kadmind</code>

**TABLE 12-1** SEAM Files (Continued)

File Name	Description
/etc/krb5/kadm5.acl	Kerberos access control list file; includes principal names of KDC administrators and their Kerberos administration privileges
/etc/krb5/kadm5.keytab	Keytab file for kadmin service on master KDC
/etc/krb5/kdc.conf	KDC configuration file
/etc/krb5/kpropd.acl	Kerberos database propagation configuration file
/etc/krb5/krb5.conf	Kerberos realm configuration file
/etc/krb5/krb5.keytab	Keytab file for network application servers
/etc/krb5/warn.conf	Kerberos warning configuration file
/etc/pam.conf	PAM configuration file
/tmp/krb5cc_uid	Default credentials cache ( <i>uid</i> is the decimal UID of the user)
/tmp/ovsec_admin.xxxxx	Temporary credentials cache for the lifetime of the password changing operation ( <i>xxxxxx</i> is a random string)
/var/krb5/.k5.REALM	KDC stash file; contains encrypted copy of the KDC master key
/var/krb5/kadmin.log	Log file for kadmin
/var/krb5/kdc.log	Log file for the KDC
/var/krb5/principal.db	Kerberos principal database
/var/krb5/principal.kadm5	Kerberos administrative database; contains policy information
/var/krb5/principal.kadm5.lock	Kerberos administrative database lock file
/var/krb5/principal.ok	Kerberos principal database initialization file; created when the Kerberos database is initialized successfully
/var/krb5/slave_datatrans	Backup file of the KDC that the <code>kprop_script</code> script uses for propagation

## PAM Configuration File

The default PAM configuration file includes entries for the authentication service, account management, session management, and password management modules.

For the authentication module, the new entries are created for `rlogin`, `login`, and `dtlogin` if SEAM 1.0 or 1.0.1 are installed. An example of these entries follows. All these services use the new PAM library, `/usr/lib/security/pam_krb5.so.1`, to provide Kerberos authentication.

These entries use the `try_first_pass` option, which requests authentication by using the user's initial password. Using the initial password means that the user is not prompted for another password, even if multiple mechanisms are listed.

```
# cat /etc/pam.conf
.
.
rlogin auth optional /usr/lib/security/pam_krb5.so.1 try_first_pass
login auth optional /usr/lib/security/pam_krb5.so.1 try_first_pass
dtlogin auth optional /usr/lib/security/pam_krb5.so.1 try_first_pass
other auth optional /usr/lib/security/pam_krb5.so.1 try_first_pass
```

For the account management module, `dtlogin` has a new entry that uses the Kerberos library, as follows. An other entry is included to provide a default rule. Currently, no actions are taken by the other entry.

```
dtlogin account optional /usr/lib/security/pam_krb5.so.1
other account optional /usr/lib/security/pam_krb5.so.1
```

The last two entries in the `/etc/pam.conf` file are shown next. The other entry for session management destroys user credentials. The new other entry for password management selects the Kerberos library.

```
other session optional /usr/lib/security/pam_krb5.so.1
other password optional /usr/lib/security/pam_krb5.so.1 try_first_pass
```

---

## SEAM Commands

This section lists some commands that are included in the SEAM product.

**TABLE 12-2** SEAM Commands

Command	Description
<code>/usr/lib/krb5/kprop</code>	Kerberos database propagation program
<code>/usr/sbin/gkadmin</code>	Kerberos database administration GUI program; used to manage principals and policies

**TABLE 12-2** SEAM Commands (Continued)

Command	Description
<code>/usr/sbin/kadmin</code>	Remote Kerberos database administration program (run with Kerberos authentication); used to manage principals, policies, and keytab files
<code>/usr/sbin/kadmin.local</code>	Local Kerberos database administration program (run without Kerberos authentication; must be run on master KDC); used to manage principals, policies, and keytab files
<code>/usr/sbin/kdb5_util</code>	Creates Kerberos databases and stash files

---

## SEAM Daemons

The following table lists the daemons that the SEAM product uses.

**TABLE 12-3** SEAM Daemons

Daemon	Description
<code>/usr/lib/krb5/kadmind</code>	Kerberos database administration daemon
<code>/usr/lib/krb5/kpropd</code>	Kerberos database propagation daemon
<code>/usr/lib/krb5/krb5kdc</code>	Kerberos ticket processing daemon

---

## SEAM Terminology

The following section presents terms and their definitions. Those terms are used throughout the SEAM documentation. In order to grasp SEAM concepts, an understanding of these terms is essential.

### Kerberos-Specific Terminology

You need to understand the terms in this section in order to administer KDCs.



The *Key Distribution Center* or *KDC* is the component of SEAM that is responsible for issuing credentials. These credentials are created by using information that is stored in the KDC database. Each realm needs at least two KDCs, a master and at least one slave. All KDCs generate credentials, but only the master KDC handles any changes to the KDC database.

A *stash file* contains an encrypted copy of the master key for the KDC. This key is used when a server is rebooted to automatically authenticate the KDC before starting the `kadmin` and `krb5kdc` commands. Because this file includes the master key, the file and any backups of the file should be kept secure. If the encryption is compromised, then the key could be used to access or modify the KDC database.

## Authentication-Specific Terminology

You need to know the terms in this section to understand the authentication process. Programmers and system administrators should be familiar with these terms.

A *client* is the software that runs on a user's workstation. The SEAM software that runs on the client makes many requests during this process. So, it is important to differentiate the actions of this software from the user.

The terms *server* and *service* are often used interchangeably. To clarify, the term *server* is used to define the physical system that SEAM software is running on. The term *service* corresponds to a particular function that is being supported on a server (for instance, `nfs`). Documentation often mentions servers as part of a service, but this definition clouds the meaning of the terms. Therefore, the term *server* refers to the physical system. The term *service* refers to the software.

The SEAM product includes three types of keys. One key is the *private key*. The private key is given to each user principal and is known only to the user of the principal and to the KDC. For user principals, the key is based on the user's password. For servers and services, the key is known as a *service key*. The service key serves the same purpose as the private key, but is used by servers and services. The third type of key is a *session key*. A session key is a key that is generated by the authentication service or the ticket-granting service. A session key is generated to provide secure transactions between a client and a service.

A *ticket* is an information packet that is used to securely pass the identity of a user to a server or service. A ticket is valid for only a single client and a particular service on a specific server. A ticket contains the principal name of the service, the principal name of the user, the IP address of the user's host, a time stamp, and a value to define the lifetime of the ticket. A ticket is created with a random session key to be used by the client and the service. After a ticket has been created, it can be reused until the ticket expires.

A *credential* is a packet of information that includes a ticket and a matching session key. Credentials are often encrypted by using either a private key or a service key, depending on which software decrypts the credential.

An *authenticator* is another type of information. When used with a ticket, an authenticator can be used to authenticate a user principal. An authenticator includes the principal name of the user, the IP address of the user's host, and a time stamp. Unlike a ticket, an authenticator can be used once only, usually when access to a service is requested. An authenticator is encrypted by using the session key for that client and that server.

## Types of Tickets

Tickets have properties that govern how they can be used. These properties are assigned to the ticket when it is created, although you can modify a ticket's properties later. For example, a ticket can change from *forwardable* to *forwarded*. You can view ticket properties with the `klist` command. See "How to View Tickets" on page 174.

Tickets can be described by one or more of the following terms:

Forwardable/forwarded	A forwardable ticket can be sent from one host to another, obviating the need for a client to reauthenticate itself. For example, if the user <code>david</code> obtains a forwardable ticket while on user <code>jennifer</code> 's machine, he can log in to his own machine without having to get a new ticket (and thus authenticate himself again). See "Example—Creating a Ticket" on page 174 for an example of a forwardable ticket.
Initial	An <i>initial</i> ticket is a ticket that is issued directly, not based on a ticket-granting ticket. Some services, such as applications that change passwords, can require tickets to be marked <i>initial</i> in order to assure themselves that the client can demonstrate a knowledge of its secret key. An <i>initial</i> ticket indicates that the client has recently authenticated itself (instead of relying on a ticket-granting ticket, which might have been around for a long time).
Invalid	An <i>invalid</i> ticket is a postdated ticket that has not yet become usable. An invalid ticket will be rejected by an application server until it becomes validated. To be validated, a ticket must be presented to the KDC by the client in a TGS request, with the <code>VALIDATE</code> flag set, after its start time has passed.
Postdatable/postdated	A <i>postdated</i> ticket is a ticket that does not become valid until some specified time after its creation. Such a ticket is useful, for example, for batch jobs that are intended to be run late at night, since the ticket, if stolen, cannot be used until the batch job is to be run. When a <i>postdated</i> ticket is issued, it is issued as <i>Invalid</i> and remains that way until:

its start time has passed, and the client requests validation by the KDC. A *postdated* ticket is normally valid until the expiration time of the ticket-granting ticket. However, if the ticket is marked *renewable*, its lifetime is normally set to be equal to the duration of the full life of the ticket-granting ticket.

#### Proxiable/proxy

At times, it is necessary for a principal to allow a service to perform an operation on its behalf. An example might be when a principal requests a service to run a print job on a third host. The service must be able to take on the identity of the client, but need only do so for that single operation. In that case, the server is said to be acting as a *proxy* for the client. The principal name of the proxy must be specified when the ticket is created.

A *proxiable* ticket is similar to a *forwardable* ticket, except that it is valid only for a single service, whereas a *forwardable* ticket grants the service the complete use of the client's identity. A *forwardable* ticket can therefore be thought of as a sort of super-proxy.

#### Renewable

Because it is a security risk to have tickets with very long lives, tickets can be designated as *renewable*. A *renewable* ticket has two expiration times: the time at which the current instance of the ticket expires, and the maximum lifetime for any ticket. If a client wants to continue to use a ticket, the client renews it before the first expiration occurs. For example, a ticket can be valid for one hour, with all tickets having a maximum lifetime of 10 hours. If the client that is holding the ticket wants to keep it for more than an hour, the client must renew it within that hour. When a ticket reaches the maximum ticket lifetime (10 hours), it automatically expires and cannot be renewed.

For information on how to view tickets to see what their attributes are, see "How to View Tickets" on page 174.

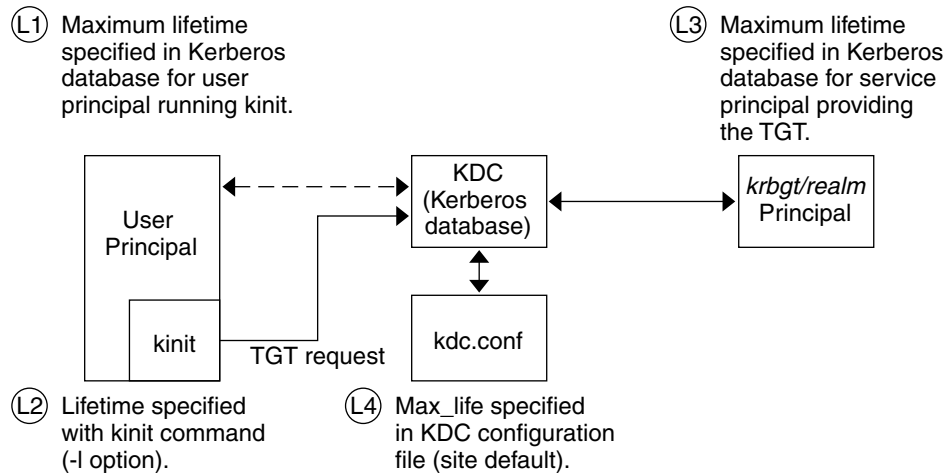
## Ticket Lifetimes

Any time a principal obtains a ticket, including a ticket-granting ticket, the ticket's lifetime is set as the smallest of the following lifetime values:

- The lifetime value that is specified by the `-l` option of `kinit`, if `kinit` is used to get the ticket.
- The maximum lifetime value (`max_life`) that is specified in the `kdc.conf` file.

- The maximum lifetime value that is specified in the Kerberos database for the service principal that provides the ticket. In the case of `kinit`, the service principal is `krbtgt/realm`.
- The maximum lifetime value that is specified in the Kerberos database for the user principal that requests the ticket.

Figure 12–1 shows how a TGT’s lifetime is determined and where the four lifetime values come from. Even though this figure shows how a TGT’s lifetime is determined, basically the same thing happens when any principal obtains a ticket. The only differences are that `kinit` doesn’t provide a lifetime value, and the service principal that provides the ticket provides a maximum lifetime value (instead of the `krbtgt/realm` principal).



Ticket lifetime = Minimum value of L1, L2, L3, and L4

**FIGURE 12–1** How a TGT’s Lifetime is Determined

The *renewable* ticket lifetime is also determined from the minimum of four values, but renewable lifetime values are used instead, as follows:

- The renewable lifetime value that is specified by the `-r` option of `kinit`, if `kinit` is used to obtain or renew the ticket.
- The maximum renewable lifetime value (`max_renewable_life`) that is specified in the `kdc.conf` file.
- The maximum lifetime renewable value that is specified in the Kerberos database for the service principal that provides the ticket. In the case of `kinit`, the service principal is `krbtgt/realm`.
- The maximum lifetime renewable value that is specified in the Kerberos database for the user principal that requests the ticket.

## Principal Names

Each ticket is identified by a principal name. The principal name can identify a user or a service. Here are examples of several principal names.

**TABLE 12–4** Examples of Principal Names

Principal Name	Description
<code>root/boston.example.com@EXAMPLE.COM</code>	A principal that is associated with the <code>root</code> account on an NFS client. This principal is called a <code>root</code> principal and is needed for authenticated NFS-mounting to succeed.
<code>host/boston.example.com@EXAMPLE.COM</code>	A principal that is used by the Kerberized applications ( <code>klist</code> and <code>kprop</code> , for example). This principal is called a <code>host</code> or <code>service</code> principal.
<code>username@EXAMPLE.COM</code>	A principal for a user.
<code>username/admin@EXAMPLE.COM</code>	An <code>admin</code> principal that can be used to administer the KDC database.
<code>nfs/boston.example.com@EXAMPLE.COM</code>	A principal that is used by the NFS service. This principal can be used instead of a <code>host</code> principal.
<code>K/M@EXAMPLE.COM</code>	The master key name principal. There is one master key name principal that is associated with each master KDC.
<code>kadmin/history@EXAMPLE.COM</code>	A principal which includes a key used to keep password histories for other principals. Each master KDC has one of these principals.
<code>kadmin/kdc1.example.com@EXAMPLE.COM</code>	A principal for the master KDC server that allows access to the KDC by using <code>kadmin</code> .
<code>changepw/kdc1.example.com@EXAMPLE.COM</code>	A principal for the master KDC server that allows access to the KDC when you are changing passwords.
<code>krbtgt/EXAMPLE.COM@EXAMPLE.COM</code>	This principal is used when you generate a ticket-granting ticket.

---

## How the Authentication System Works

Applications allow you to log in to a remote system if you can provide a ticket that proves your identity, and a matching session key. The session key contains information that is specific to the user and the service that is being accessed. A ticket and session key are created by the KDC for all users when they first log in. The ticket

and the matching session key form a credential. While using multiple networking services, a user can gather many credentials. The user needs to have a credential for each service that runs on a particular server. For instance, access to the `ftp` service on a server named `boston` requires one credential. Access to the `ftp` service on another server requires its own credential.

The process of creating and storing the credentials is transparent. Credentials are created by the KDC that sends the credential to the requester. When received, the credential is stored in a credential cache.

---

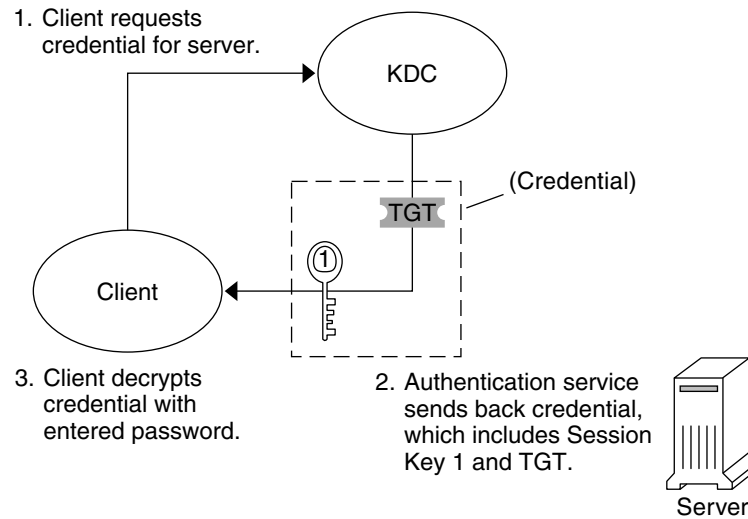
## Gaining Access to a Service Using SEAM

In order for a user to access a specific service on a specific server, the user must obtain two credentials. The first credential is for the ticket-granting service (known as the TGT). Once the ticket-granting service has decrypted this credential, the service creates a second credential for the server that the user is requesting access to. This second credential can then be used to request access to the service on the server. After the server has successfully decrypted the second credential, then the user is given access. The following sections describe this process in more detail.

### Obtaining a Credential for the Ticket-Granting Service

1. To start the authentication process, the client sends a request to the authentication server for a specific user principal. This request is sent without encryption. There is no secure information included in the request, so it is not necessary to use encryption.
2. When the request is received by the authentication service, the principal name of the user is looked up in the KDC database. If a principal matches, the authentication service obtains the private key for that principal. The authentication service then generates a session key to be used by the client and the ticket-granting service (call it session key 1) and a ticket for the ticket-granting service (ticket 1). This ticket is also known as the ticket-granting ticket (TGT). Both the session key and the ticket are encrypted by using the user's private key, and the information is sent back to the client.
3. The client uses this information to decrypt session key 1 and ticket 1, by using the private key for the user principal. Since the private key should only be known by the user and the KDC database, the information in the packet should be safe. The client stores the information in the credentials cache.

During this process, a user is normally prompted for a password. If the password the user enters is the same as the password that was used to build the private key stored in the KDC database, then the client can successfully decrypt the information that is sent by the authentication service. Now the client has a credential to be used with the ticket-granting service. The client is ready to request a credential for a server.



TGT = Ticket-granting ticket  
KDC = Key Distribution Center

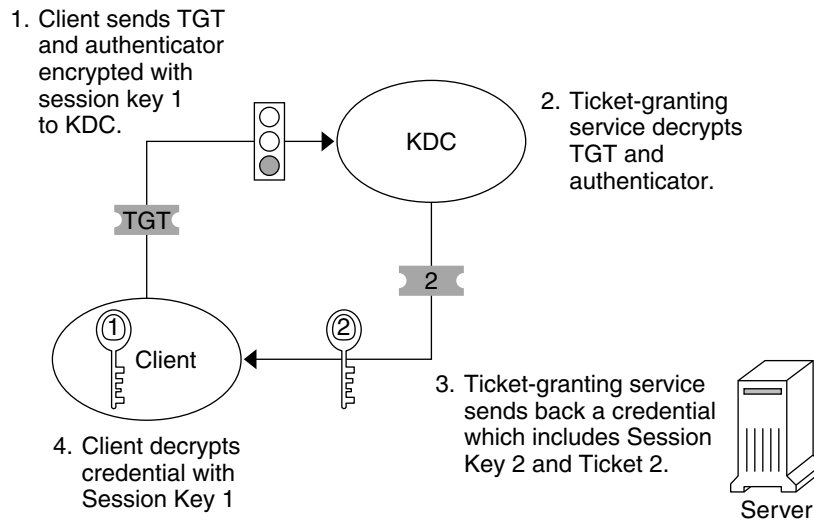
FIGURE 12-2 Obtaining a Credential for the Ticket-Granting Service

## Obtaining a Credential for a Server

1. To request access to a specific server, a client must first have obtained a credential for that server from the authentication service. See “Obtaining a Credential for the Ticket-Granting Service” on page 190. The client then sends a request to the ticket-granting service, which includes the service principal name, ticket 1, and an authenticator that was encrypted with session key 1. Ticket 1 was originally encrypted by the authentication service by using the service key of the ticket-granting service.
2. Because the service key of the ticket-granting service is known to the ticket-granting service, ticket 1 can be decrypted. The information in ticket 1 includes session key 1, so the ticket-granting service can decrypt the authenticator. At this point, the user principal is authenticated with the ticket-granting service.
3. Once the authentication is successful, the ticket-granting service generates a session key for the user principal and the server (session key 2), and a ticket for the server (ticket 2). Session key 2 and ticket 2 are then encrypted by using session key 1. Since session key 1 is known only to the client and the ticket-granting service, this

information is secure and can be safely sent over the network.

4. When the client receives this information packet, the client decrypts the information by using session key 1, which it had stored in the credential cache. The client has obtained a credential to be used with the server. Now the client is ready to request access to a particular service on that server.

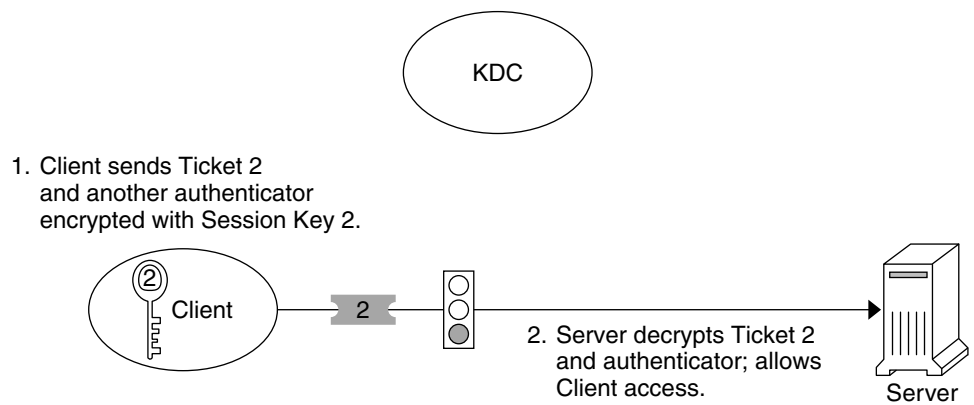


**FIGURE 12-3** Obtaining a Credential for a Server

## Obtaining Access to a Specific Service

1. To request access to a specific service, the client must first have obtained a credential for the ticket-granting service from the authentication server, and a server credential from the ticket-granting service. See “Obtaining a Credential for the Ticket-Granting Service” on page 190 and “Obtaining a Credential for a Server” on page 191. The client can send a request to the server including ticket 2 and another authenticator. The authenticator is encrypted by using session key 2.
2. Ticket 2 was encrypted by the ticket-granting service with the service key for the service. Since the service key is known by the service principal, the service can decrypt ticket 2 and get session key 2. Session key 2 can then be used to decrypt the authenticator. If the authenticator is successfully decrypted, the client is given access to the service.





**FIGURE 12-4** Obtaining Access to a Specific Service

## Using the `gsscred` Table

The `gsscred` table is used by an NFS server when the server is trying to identify a SEAM user. The NFS services use UNIX IDs to identify users. These IDs are not part of a user principal or credential. The `gsscred` table provides a mapping from UNIX UIDs (from the password file) to principal names. The table must be created and administered after the KDC database is populated.

When a client request comes in, the NFS services try to map the principal name to a UNIX ID. If the mapping fails, the `gsscred` table is consulted. With the `kerberos_v5` mechanism, a `root/hostname` principal is automatically mapped to UID 0, and the `gsscred` table is not consulted. Thus, there is no way to do special remappings of `root` through the `gsscred` table.



## Managing System Security Topics

---

This topic map identifies the chapters that are related to managing system security in the Solaris environment.

Chapter 14	Provides overview information about file security, system security, and network security.
Chapter 15	Provides step-by-step instructions on how to display file information, change file ownership and file permissions, and set special file permissions.
Chapter 16	Provides step-by-step instructions on how to check login status, set up dial-up passwords, restrict root access, and monitor root access and <code>su</code> attempts.
Chapter 17	Provides overview information for using role-based access control (RBAC).
Chapter 18	Provides step-by-step instructions for using RBAC.
Chapter 19	Provides reference information about RBAC.
Chapter 20	Provides overview information about the Automated Security Enhancement Tool (ASET) and step-by-step instructions on how to run ASET interactively or periodically (by using a <code>cron</code> job). This chapter also includes information about collecting client ASET reports on a server.



## Managing System Security (Overview)

---

Keeping a system's information secure is an important system administration responsibility. This chapter provides overview information about managing system security at the file, system, and network level.

This is a list of the overview information in this chapter.

- "Controlling Access to a Computer System" on page 197
- "File Security" on page 200
- "System Security" on page 202
- "Network Security" on page 206

---

## Controlling Access to a Computer System

At the file level, the SunOS operating system provides some standard security features that you can use to protect files, directories, and devices. At the system and network levels, the security issues are mostly the same. In the workplace, a number of systems that are connected to a server can be thought of as one large multifaceted system. The system administrator is responsible for the security of this larger system or network. Not only is it important to defend the network from outsiders who are trying to gain access to the network, but it is also important to ensure the integrity of the data on the systems within the network.

The first line of security defense is to control access to your system. You can control and monitor system access by doing the following:

- Maintaining physical site security
- Maintaining login control
- Restricting access to data in files

- Maintaining network control
- Monitoring system usage
- Setting the path variable correctly
- Securing files
- Installing a firewall
- Reporting security problems

## Maintaining Physical Site Security

To control access to your system, you must maintain the physical security of your computing environment. For instance, if a system is logged in and left unattended, anyone who can use that system can gain access to the operating system and the network. You need to be aware of your computer's surroundings and physically protect it from unauthorized access.

## Maintaining Login Control

You also must restrict unauthorized logins to a system or the network, which you can do through password and login control. All accounts on a system should have a password. An account without a password makes your entire network accessible to anyone who can guess a user name.

Solaris software restricts control of certain system devices to the user login account. Only a process that is running as superuser or console user can access a system mouse, keyboard, frame buffer, or audio device unless the `/etc/logindevperm` file is edited. For more information, see `logindevperm(4)`.

## Restricting Access to Data in Files

After you have established login restrictions, you can control access to the data on your system. You might want to allow some users to read some files, and give other users permission to change or delete some files. You might have some data that you do not want anyone else to see. Chapter 15 discusses how to set file permissions.

## Maintaining Network Control

Computers are often part of a configuration of systems called a *network*. A network allows connected systems to exchange information and access data and other resources that are available from systems connected to the network. Networking has created a powerful and sophisticated way of computing. However, networking has also jeopardized computer security.

For instance, within a network of computers, individual systems are open to allow sharing of information. Also, because many people have access to the network, there is more chance for allowing unwanted access, especially through user error (for example, through a poor use of passwords).

## Monitoring System Usage

As system administrator, you need to monitor system activity, being aware of all aspects of your systems, including the following:

- What is the normal load?
- Who has access to the system?
- When do individuals access the system?

With this kind of knowledge, you can use the available tools to audit system use and monitor the activities of individual users. Monitoring is very useful when there is a suspected breach in security.

## Setting the Correct Path

It is important to set your path variable correctly. Otherwise, you can accidentally run a program that was introduced by someone else that harms your data or your system. This kind of program, which creates a security hazard, is referred to as a “Trojan horse.” For example, a substitute `su` program could be placed in a public directory where you, as system administrator, might run it. Such a script would look just like the regular `su` command. Since the script removes itself after execution, it is hard to tell that you have actually run a Trojan horse.

The path variable is automatically set at login time through the startup files: `.login`, `.profile`, and `.cshrc`. Setting up the user search path so that the current directory (`.`) comes last prevents you or your users from running this type of Trojan horse. The path variable for superuser should not include the current directory at all. The Automated Security Enhancement Tool (ASET) examines the startup files to ensure that the path variable is set up correctly and that it does not contain a dot (`.`) entry.

## Securing Files

Since the SunOS operating system is a multiuser system, file system security is the most basic, and most important, security risk on a system. You can use both the traditional UNIX file protection or the more secure access control lists (ACLs) to protect your files.

Also, many executable programs have to be run as `root` (that is, as superuser) to work properly. These executables run with the user ID set to 0 (`setuid=0`). Anyone who is running these programs runs them with the `root` ID, which creates a potential security problem if the programs are not written with security in mind.

Except for the executables that are shipped with the `setuid` bit set to `root`, you should disallow the use of `setuid` programs, or at least restrict them and keep them to a minimum.

## Installing a Firewall

Another way to protect your network is to use a firewall or secure gateway system. A firewall is a dedicated system that separates two networks, each of which approaches the other as untrusted. You should consider this setup as mandatory between your internal network and any external networks, such as the Internet, with which you want internal network users to communicate.

A firewall can also be useful between some internal networks. For example, the firewall or secure gateway computer will not send a packet between two networks unless the gateway computer is the origin or the destination address of the packet. A firewall should also be set up to forward packets for particular protocols only. For example, you can allow packets for transferring mail, but not those packets for the `telnet` or `rlogin` command. ASET, when run at high security, disables the forwarding of Internet Protocol (IP) packets.

## Reporting Security Problems

If you experience a suspected security breach, you can contact the Computer Emergency Response Team/Coordination Center (CERT/CC). CERT/CC is a Defense Advanced Research Projects Agency (DARPA) funded project that is located at the Software Engineering Institute at Carnegie Mellon University. This agency can assist you with any security problems you are having. This agency can also direct you to other Computer Emergency Response Teams that might be more appropriate for your particular needs. You can call CERT/CC at its 24-hour hotline: (412) 268-7090, or contact the team by email at `cert@cert.sei.cmu.edu`.

---

## File Security

The SunOS operating system is a multiuser system, which means that all the users who are logged in to a system can read and use files that belong to one another, as long as they have the file permissions to do so. Table 14-1 describes the commands for file system security. For step-by-step instructions on securing files, see Chapter 15.

## Commands for File System Security

This table describes the commands for monitoring and securing files and directories.



**TABLE 14-1** Commands for File System Security

Command	Description	Man Page
ls	Lists the files in a directory and information about them.	ls(1)
chown	Changes the ownership of a file.	chown(1)
chgrp	Changes the group ownership of a file.	chgrp(1)
chmod	Changes permissions on a file. You can use either symbolic mode (letters and symbols) or absolute mode (octal numbers) to change permissions on a file.	chmod(1)

## File Encryption

By placing a sensitive file into an inaccessible directory (700 mode) and making the file unreadable by other users (600 mode), you will keep it secure in most cases. However, someone who guesses your password or the `root` password can read and write to that file. Also, the sensitive file is preserved on a backup tape every time you back up the system files to tape.

Fortunately, an additional layer of security is available to all SunOS system software users in the United States: the optional file encryption kit. The encryption kit includes the `crypt` command, which scrambles the data to disguise the text. For more information, see `crypt(1)`.

## Access Control Lists (ACLs)

ACLs (pronounced “ackkls”) can provide greater control over file permissions when the traditional UNIX file protection in the SunOS operating system is not sufficient. The traditional UNIX file protection provides read, write, and execute permissions for the three user classes: owner, group, and other. An ACL provides better file security by enabling you to define file permissions for the owner, owner’s group, others, specific users and groups, and to define default permissions for each of those categories. For step-by-step instructions on using ACLs, see “Using Access Control Lists (ACLs)” on page 226.

The following table lists the commands for administering ACLs on files or directories.

TABLE 14-2 ACL Commands

Command	Description	Man Page
setfacl	Sets, adds, modifies, and deletes ACL entries	setfacl(1)
getfacl	Displays ACL entries	getfacl(1)

---

## System Security

This section describes how to safeguard your system against unauthorized access, such as the following:

- How to prevent an intruder from logging in to your system
- How to maintain the password files
- How to prevent unauthorized superuser access to sensitive system files and programs

You can set up two security barriers on a system. The first security barrier is the `login` command. To cross this barrier and gain access to a system, a user must supply a user name and a corresponding password that is known by the local system or by the name service (LDAP, NIS, or NIS+).

The second security barrier is ensuring that the system files and programs can be changed or removed by superuser only. A would-be superuser must supply the `root` user name and its correct password.

## Login Access Restrictions

When a user logs in to a system, the `login` command consults the appropriate database according to the information that is listed in the `/etc/nsswitch.conf` file. The entries in this file can include `files` (which designates the `/etc` files), `nis` (which designates the NIS database), `ldap` (which designates the LDAP directory service), and `nisplus` (which designates the NIS+ database). For a description of this file, see the `nsswitch.conf(4)` man page. For information about naming or directory services, see *System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)* or *System Administration Guide: Naming and Directory Services (NIS and NIS+)*.

The `login` command verifies the user name and password that were entered. If the user name is not in the password file or the password is not correct for the user name, the `login` command denies access to the system. When the user supplies a user name from the password file and the correct password for the user name, the system grants the user access to the system.

## Special Logins

There are two common ways to access a system: by using a conventional user login, or by using the `root` login. In addition, a number of special *system* logins allow a user to perform administrative commands without using the `root` account. The administrator assigns passwords to these login accounts.

The following table lists some system login accounts and their uses. The system logins perform special functions, and each login has its own group identifier number (GID). Each of these logins should have its own password, which should be distributed on a need-to-know basis.

**TABLE 14-3** System Logins

Login Account	GID	Use
<code>root</code>	0	Has almost no restrictions and overrides all other logins, protections, and permissions. The <code>root</code> account has access to the entire system. The password for the <code>root</code> login should be very carefully protected. The <code>root</code> account owns most of the Solaris commands.
<code>daemon</code>	1	Controls background processing.
<code>bin</code>	2	Owns some of the Solaris commands.
<code>sys</code>	3	Owns many system files.
<code>adm</code>	4	Owns certain administrative files.
<code>lp</code>	71	Owns the object data files and spooled data files for the printer.
<code>uucp</code>	5	Owns the object data files and spooled data files for UUCP, the UNIX-to-UNIX copy program.
<code>nuucp</code>	9	Is used by remote systems to log in to the system and start file transfers.

You should also set the security of the `eeeprom` command to require a password. See `eeeprom(1M)` for more information.

## Managing Password Information

When users log in to a system, they must enter both a user name and a password. Although logins are publicly known, passwords must be kept secret, known only to each user. You should ask your users to choose their passwords carefully, and they should change them often.

Passwords are initially created when you set up a user account. To maintain security on user accounts, you can set up password aging to force users to routinely change their passwords. You can also disable a user account by locking the password. For detailed information about setting up and maintaining passwords, see “Managing User Accounts and Groups (Overview)” in *System Administration Guide: Basic Administration* and the `passwd(1)` man page.

## NIS+ Password File

If your network uses NIS+, the password information is kept in the NIS+ database. Information in the NIS+ database can be protected by restricting access to authorized users. You can use the `passwd` command to change a user’s NIS+ password.

## NIS Password File

If your network uses NIS, the password information is kept in the NIS password map. NIS does not support password aging. You can use the `passwd` command to change a user’s NIS password.

## /etc Files

If your network uses `/etc` files, the password information is kept in the system’s `/etc/passwd` and `/etc/shadow` files. The user name and other information are kept in the password file `/etc/passwd`, while the encrypted password itself is kept in a separate *shadow* file, `/etc/shadow`. This security measure prevents a user from gaining access to the encrypted passwords. While the `/etc/passwd` file is available to anyone who can log in to a machine, only superuser can read the `/etc/shadow` file. You can use the `passwd` command to change a user’s password on a local system.

## LDAP Password Information

If your network uses LDAP, the password and shadow information is stored in the `ou=people` container of the LDAP directory tree. You can use the `password -r ldap` command to change a user’s LDAP password.

## Using the Restricted Shell

The standard shell allows a user to open files, execute commands, and so on. The restricted shell (`rsh`) can be used to limit the ability of a user to change directories and execute commands. The restricted shell is located in the `/usr/lib` directory. Note that the restricted shell is not the remote shell, which is `/usr/sbin/rsh`. The restricted shell differs from the standard shell in these ways:

- The user is limited to the home directory, so the user cannot use the `cd` command to change directories.
- The user can use only commands in the `PATH` set by the system administrator, so the user cannot change the `PATH` variable.
- The user can access only files in the home directory and its subdirectories, so the user cannot name commands or files by using a complete path name.
- The user cannot redirect output with `>` or `>>`.

The restricted shell allows the system administrator to limit a user's ability to stray into the system files, and is intended mainly to set up a user who needs to perform specific tasks. The restricted shell is not completely secure, however, and is only intended to keep unskilled users from getting into (or causing) trouble.

For information about the restricted shell, see the `rsh(1M)` man page.

A more secure alternative to the restricted shell is the Secure Shell (`ssh`), which enables users to securely access a remote host over an unsecured network. For information about using the Secure Shell, see Chapter 5.

## Tracking Superuser (Root) Login

Your system requires a `root` password for superuser mode. In the default configuration, a user cannot remotely log in to a system as `root`. When logging in remotely, a user must log in as himself or herself and then use the `su` command to become `root`. This setup enables you to track who is using superuser privileges on your system.

## Monitoring Who is Becoming Superuser or Other Users

You have to use the `su` command to change to another user, for example, if you want to become superuser. For security reasons, you can monitor who has been using the `su` command, especially those users who are trying to gain superuser access.

For detailed instructions, see "How to Monitor Who Is Using the `su` Command" on page 243.

---

# Network Security

The more available access is across a network, the more advantageous it is for networked systems. However, free access and the sharing of data and resources create security problems. Network security is usually based on limiting or blocking operations from remote systems. The following figure describes the security restrictions that you can impose on remote operations.

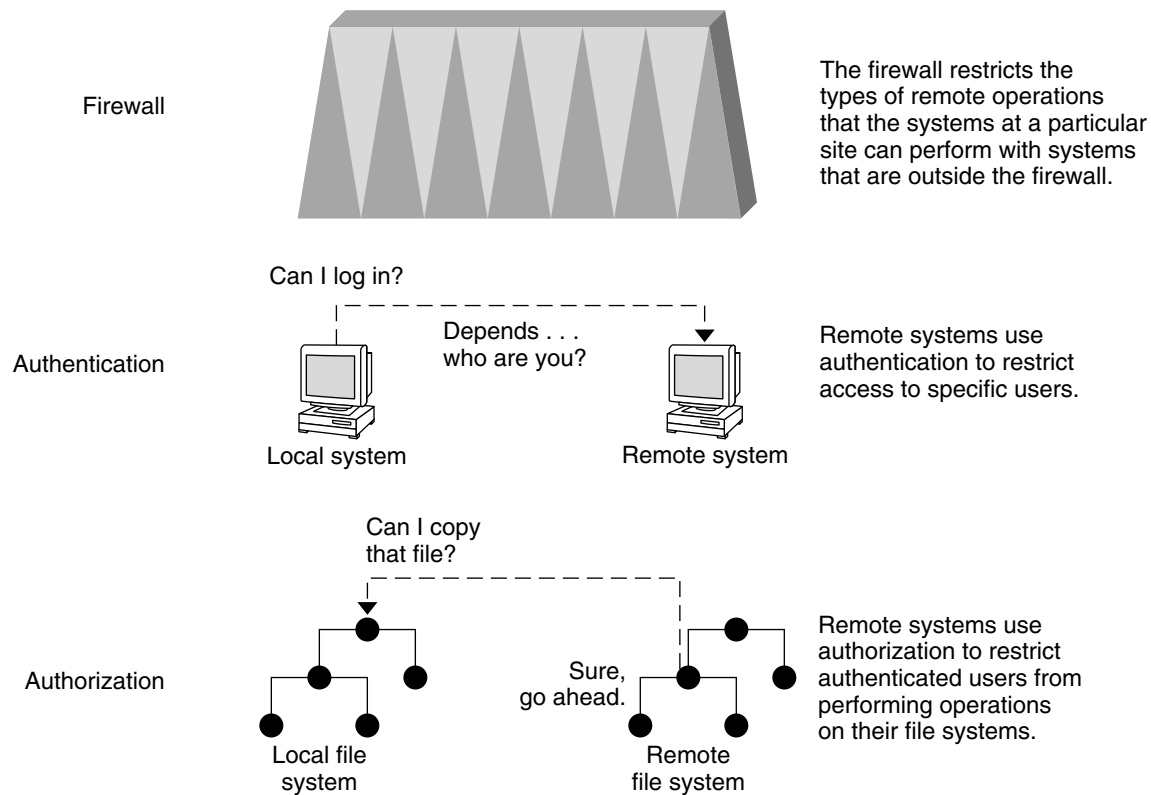


FIGURE 14-1 Security Restrictions for Remote Operations

## Firewall Systems

You can set up a firewall system to protect the resources in your network from outside access. A *firewall system* is a secure host that acts as a barrier between your internal network and outside networks.

The firewall has two functions. The firewall acts as a gateway that passes data between the networks, and it acts as a barrier that blocks the free passage of data to and from the network. The firewall requires a user on the internal network to log in to the firewall system to access hosts on remote networks. Similarly, a user on an outside network must log in to the firewall system before being granted access to a host on the internal network.

In addition, all electronic mail that is sent from the internal network is sent to the firewall system for transfer to a host on an external network. The firewall system receives all incoming electronic mail, and distributes it to the hosts on the internal network.



---

**Caution** – A firewall prevents unauthorized users from accessing the hosts on your network. You should maintain strict and rigidly enforced security on the firewall, but security on other hosts on the network can be more relaxed. However, an intruder who can break into your firewall system can then gain access to all the other hosts on the internal network.

---

A firewall system should not have any *trusted hosts*. A trusted host is a host from which a user can log in without being required to type in a password. A firewall system should not share any of its file systems, or mount any file systems from other servers.

ASET can be used to make a system into a firewall, and to enforce high security on a firewall system, as described in Chapter 20.

## Packet Smashing

Most local area networks transmit data between computers in blocks called packets. Through a procedure called *packet smashing*, unauthorized users can harm or destroy data. Packet smashing involves capturing the packets before they reach their destination, injecting arbitrary data into the contents, and then sending the packets back on their original course. On a local area network, packet smashing is impossible because packets reach all systems, including the server, at the same time. Packet smashing is possible on a gateway, however, so make sure that all gateways on the network are protected.

The most dangerous attacks are those attacks that affect the integrity of the data. Such attacks involve changing the contents of the packets or impersonating a user. Attacks that involve eavesdropping—recording conversations and replaying them later without impersonating a user—do not compromise data integrity. These attacks do affect privacy, however. You can protect the privacy of sensitive information by encrypting data that goes over the network.

## Authentication and Authorization

*Authentication* is a way to restrict access to specific users when they access a remote system, which can be set up at both the system level or network level. Once a user gains access to a remote system, *authorization* is a way to restrict operations that the user can perform on the remote system. The following table lists the types of authentications and authorizations that can help protect your systems on the network against unauthorized use.

**TABLE 14-4** Types of Authentication and Authorization for Remote Access

Type	Description	Where to Find Information
LDAP and NIS+	The LDAP directory service and the NIS+ name service can provide both authentication and authorization at the network level.	<i>System Administration Guide: Naming and Directory Services (DNS, NIS, and LDAP)</i> and <i>System Administration Guide: Naming and Directory Services (FNS and NIS+)</i>
Remote login commands	The remote login commands ( <code>rlogin</code> , <code>rcp</code> , <code>ftp</code> ) enable users to log in to a remote system over the network and use its resources. If you are a "trusted host," authentication is automatic. Otherwise, you are asked to authenticate yourself.	"Accessing Remote Systems (Tasks)" in <i>System Administration Guide: Resource Management and Network Services</i>
Secure RPC	Secure RPC improves the security of network environments by authenticating users who make requests on remote systems. You can use either the UNIX, DES, or Kerberos authentication system for Secure RPC.	"Overview of Secure RPC" on page 27
	Secure RPC can also be used to provide additional security to the NFS environment, called Secure NFS.	"NFS Services and Secure RPC" on page 27
DES encryption	The Data Encryption Standard (DES) encryption functions use a 56-bit key to encrypt a secret key.	"DES Encryption" on page 28
Diffie-Hellman authentication	This authentication method is based on the ability of the sending system to use the common key to encrypt the current time, which the receiving system can decrypt and check against its current time.	"Diffie-Hellman Authentication" on page 28
Kerberos	Kerberos uses DES encryption to authenticate a user when logging in to the system.	Chapter 3



## Sharing Files

A network file server can control which files are available for sharing. A network file server can also control which clients have access to the files, and what type of access is permitted for those clients. In general, the file server can grant read and write access or read-only access either to all clients or to specific clients. Access control is specified when resources are made available with the `share` command.

A server can use the `/etc/dfs/dfstab` file to list the file systems that it makes available to clients on the network. For more information about sharing file systems, see “Automatic File-System Sharing” in *System Administration Guide: Resource Management and Network Services*.

## Restricting Superuser (Root) Access

In general, superuser is not allowed `root` access to file systems that are shared across the network. Unless the server specifically grants superuser privileges, a user who is logged in as superuser on a client cannot gain `root` access to files that are remotely mounted on the client. The NFS system implements this strategy by changing the user of the requester to the user, `nobody` (user ID 60001). The access rights of user `nobody` are the same as those access rights that are given to the public or a user without credentials. For example, if the public has only execute permission for a file, then user `nobody` can only execute that file.

An NFS server can grant superuser privileges on a shared file system on a per-host basis by using the `root=hostname` option to the `share` command.

## Using Privileged Ports

If you do not want to run Secure RPC, a possible substitute is the Solaris “privileged port” mechanism. A privileged port is assigned with a port number of less than 1024. After a client system has authenticated the client’s credential, it builds a connection to the server by using the privileged port. The server then verifies the client credential by examining the connection’s port number.

Non-Solaris clients, however, might be unable to communicate by using the privileged port. If they cannot, you see an error message similar to the following:

```
“Weak Authentication  
NFS request from unprivileged port”
```

## Using the Automated Security Enhancement Tool (ASET)

The ASET security package provides automated administration tools that enable you to control and monitor your system's security. You specify a security level (low, medium, or high) at which ASET will run. At each higher level, ASET's file-control functions increase to reduce file access and tighten your system security.

For more information, see Chapter 20.

## Securing Files (Tasks)

---

This chapter describes the procedures for securing files in the Solaris environment.

This is a list of the step-by-step instructions in this chapter.

- “How to Display File Information” on page 215
- “How to Change the Owner of a File” on page 217
- “How to Change Group Ownership of a File” on page 218
- “How to Change Permissions in Absolute Mode” on page 221
- “How to Change Special Permissions in Absolute Mode” on page 222
- “How to Change Permissions in Symbolic Mode” on page 223
- “How to Find Files With `setuid` Permissions” on page 224
- “How to Disable Programs From Using Executable Stacks” on page 226
- “How to Disable Executable Stack Message Logging” on page 226
- “How to Set an ACL on a File” on page 229
- “How to Copy an ACL” on page 231
- “How to Check If a File Has an ACL” on page 231
- “How to Modify ACL Entries on a File” on page 232
- “How to Delete ACL Entries From a File” on page 233
- “How to Display ACL Entries for a File” on page 233

---

## File Security Features

This section describes the features that constitute a file’s security.

### User Classes

For each file, there are three classes of users that specify the levels of security:

- user – The file or directory owner, which is usually the user who created the file. The owner of a file can decide who has the right to read it, to write to it (make changes to it), or, if it is a command, to execute it.
- group – Members of a group.
- others – All other users who are not the file owner or group owner.

Only the owner of the file or `root` can assign or modify file permissions.

## File Permissions

The following table lists and describes the permissions that you can give to each user class for a file.

**TABLE 15–1** File Permissions

Symbol	Permission	Description
r	Read	Designated users can open and read the contents of a file.
w	Write	Designated users can write to the file (modify its contents), add to it, or delete it.
x	Execute	Designated users can execute the file (if it is a program or shell script), or run the program with one of the <code>exec (2)</code> system calls.
-	Denied	Designated users cannot read, write, or execute the file.

These file permissions apply to special files such as devices, sockets, and named pipes (FIFOs), as they do to regular files.

For a symbolic link, the permissions that apply are the permissions of the file that the link points to.

## Directory Permissions

The following table lists and describes the permissions that you can give to each user class for a directory.

**TABLE 15–2** Directory Permissions

Symbol	Permission	Description
r	Read	Designated users can list files in the directory.

**TABLE 15-2** Directory Permissions (Continued)

Symbol	Permission	Description
w	Write	Designated users can add or remove files or links in the directory.
x	Execute	Designated users can open or execute files in the directory. They also can make the directory and the directories beneath it current.

You can protect the files in a directory (and in its subdirectories) by disallowing access to that directory by setting restrictive file permissions. Note, however, that superuser has access to all files and directories on the system.

## Special File Permissions (setuid, setgid and Sticky Bit)

Three special types of permissions are available for executable files and public directories. When these permissions are set, any user who runs that executable file assumes the user ID of the owner (or group) of the executable file.

You must be extremely careful when you set special permissions, because special permissions constitute a security risk. For example, a user can gain superuser privileges by executing a program that sets the user ID (UID) to `root`. Also, all users can set special permissions for files they own, which constitutes another security concern.

You should monitor your system for any unauthorized use of the `setuid` and `setgid` permissions to gain superuser privileges. To search for and list all of the files that use these permissions, see “How to Find Files With `setuid` Permissions” on page 224. A suspicious listing grants ownership of such a program to a user rather than to `root` or `bin`.

### setuid Permission

When set-user identification (`setuid`) permission is set on an executable file, a process that runs this file is granted access based on the owner of the file (usually `root`), rather than the user who is running the executable file. This special permission allows a user to access files and directories that are normally only available to the owner. For example, the `setuid` permission on the `passwd` command makes it possible for a user to change passwords, assuming the permissions of the `root` ID:

```
-r-sr-sr-x  3 root  sys      104580 Sep 16 12:02 /usr/bin/passwd
```

This special permission presents a security risk, because some determined users can find a way to maintain the permissions that are granted to them by the `setuid` process even after the process has finished executing.

---

**Note** – The use of `setuid` permissions with the reserved UIDs (0–100) from a program might not set the effective UID correctly. Use a shell script instead or avoid using the reserved UIDs with `setuid` permissions.

---

## setgid Permission

The set-group identification (`setgid`) permission is similar to `setuid`, except that the process's effective group ID (GID) is changed to the group owner of the file, and a user is granted access based on permissions granted to that group. The `/usr/bin/mail` command has `setgid` permissions:

```
-r-x--s--x  1 root    mail      63628 Sep 16 12:01 /usr/bin/mail
```

When `setgid` permission is applied to a directory, files that were created in this directory belong to the group to which the directory belongs, not the group to which the creating process belongs. Any user who has write and execute permissions in the directory can create a file there. However, the file belongs to the group that owns the directory, not to the user's group ownership.

You should monitor your system for any unauthorized use of the `setuid` and `setgid` permissions to gain superuser privileges. To search for and list all of the files that use these permissions, see "How to Find Files With `setuid` Permissions" on page 224. A suspicious listing grants group ownership of such a program to a user rather than to `root` or `bin`.

## Sticky Bit

The *sticky bit* is a permission bit that protects the files within a directory. If the directory has the sticky bit set, a file can be deleted only by the owner of the file, the owner of the directory, or by `root`. This special permission prevents a user from deleting other users' files from public directories such as `/tmp`:

```
drwxrwxrwt 7  root  sys   400 Sep  3 13:37 tmp
```

Be sure to set the sticky bit manually when you set up a public directory on a TMPFS file system.

## Default umask Setting

When you create a file or directory, it has a default set of permissions. These default permissions are determined by the `umask` setting in the `/etc/profile` file, or in your `.cshrc` or `.login` file. By default, the system sets the permissions on a text file to 666, which grants read and write permission to user, group, and others, and to 777 on a directory or executable file.

The value assigned by the `umask` command is subtracted from the default. This process has the effect of denying permissions in the same way that the `chmod` command grants them. For example, while the `chmod 022` command grants write permission to group and others, the `umask 022` command denies write permission for group and others.

The following table shows some typical `umask` settings, and the effect on an executable file.

**TABLE 15-3** `umask` Settings for Different Security Levels

Level of Security	<code>umask</code> Setting	Permissions Disallowed
Permissive (744)	022	w for group and others
Moderate (740)	027	w for group, <code>rx</code> for others
Moderate (741)	026	w for group, <code>r</code> for others
Severe (700)	077	<code>rx</code> for group and others

For more information on setting the `umask` value, see the `umask(1)` man page.

---

## Displaying File Information

This section describes how to display file information.

### ▼ How to Display File Information

Display information about all the files in a directory by using the `ls` command.

```
$ ls -la
```

- l Displays the long format that includes user and group ownership and file permissions.
- a Displays all files, including hidden files that begin with a dot (.).

Each line in the display has the following information about a file:

- Type of file

A file can be one of seven types. The following table lists the possible file types.

**TABLE 15-4** File Types

Symbol	Type
-	Text or program
D	Door
d	Directory
b	Block special file
c	Character special file
p	Named pipe (FIFO)
l	Symbolic link
s	Socket

- Permissions; see Table 15-1 and Table 15-2 for descriptions
- Number of hard links
- Owner of the file
- Group of the file
- Size of the file, in bytes
- Date the file was created or the last date that the file was changed
- Name of the file

### Example—Displaying File Information

The following example displays the partial list of the files in the /sbin directory.

```
$ cd /sbin
$ ls -la
total 13456
drwxr-xr-x  2 root    sys      512 Sep  1 14:11 .
drwxr-xr-x 29 root    root     1024 Sep  1 15:40 ..
-r-xr-xr-x  1 root    bin     218188 Aug 18 15:17 autopush
lrwxrwxrwx  1 root    root     21 Sep  1 14:11 bpgetfile -> ...
-r-xr-xr-x  1 root    bin    505556 Aug 20 13:24 dhcpagent
-r-xr-xr-x  1 root    bin    456064 Aug 20 13:25 dhcpinfo
-r-xr-xr-x  1 root    bin    272360 Aug 18 15:19 fdisk
-r-xr-xr-x  1 root    bin    824728 Aug 20 13:29 hostconfig
-r-xr-xr-x  1 root    bin    603528 Aug 20 13:21 ifconfig
-r-xr-xr-x  1 root    sys    556008 Aug 20 13:21 init
-r-xr-xr-x  2 root    root    274020 Aug 18 15:28 jsh
-r-xr-xr-x  1 root    bin    238736 Aug 21 19:46 mount
-r-xr-xr-x  1 root    sys     7696 Aug 18 15:20 mountall
.
.
.
```



---

## Changing File Ownership

This section describes how to change the ownership and group ownership of a file.

By default, the owner cannot use the `chown` command to change the owner of a file or directory. However, you can enable the owner to use the `chown` command by adding the following line to the system's `/etc/system` file and rebooting the system.

```
set rstchown = 0
```

For more information, see `chown(1)`.

In addition, the owner can only use the `chgrp` command to change the group of a file to a group in which the owner belongs by default. For example, if the owner of a file only belongs to the `staff` and `sysadm` groups, the owner can only change the group of a file to `staff` or `sysadm` group.

However, you can enable the owner to change the group of a file to a group in which the owner doesn't belong by adding the following line to the system's `/etc/system` file and rebooting the system.

```
set rstchown = 0
```

For more information, see `chgrp(1)`.

Also, be aware that there can be other restrictions on changing ownership and groups on NFS-mounted file systems.

### ▼ How to Change the Owner of a File

Use the following procedure to change the ownership of a file.

1. **Become superuser or assume an equivalent role.**
2. **Change the owner of a file by using the `chown` command.**

```
# chown new-owner filename
```

*new-owner* Specifies the user name or UID of the new owner of the file or directory.

*filename* Specifies the file or directory.

3. **Verify that the owner of the file has changed.**

```
# ls -l filename
```

## Example—Changing the Owner of a File

In the following example, the ownership on `myfile` is changed to the user `rimmer`.

```
# chown rimmer myfile
# ls -l myfile
-rw-r--r--  1 rimmer  scifi  112640 May 24 10:49 myfile
```

## ▼ How to Change Group Ownership of a File

Use the following procedure to change the group ownership of a file.

1. **Become superuser or assume an equivalent role.**
2. **Change the group owner of a file by using the `chgrp` command.**

```
$ chgrp group filename
```

<i>group</i>	Specifies the group name or GID of the new group of the file or directory.
<i>filename</i>	Specifies the file or directory.

For information on setting up groups, see “Managing User Accounts and Groups (Overview)” in *System Administration Guide: Basic Administration*.

3. **Verify that the group owner of the file has changed.**

```
$ ls -l filename
```

## Example—Changing Group Ownership of a File

In the following example, the group ownership on `myfile` is changed to the group `scifi`.

```
$ chgrp scifi myfile
$ ls -l myfile
-rwxrw--  1 rimmer  scifi  12985 Nov 12 16:28 myfile
```

---

## Changing File Permissions

The `chmod` command enables you to change the permissions on a file. You must be superuser or the owner of a file or directory to change its permissions.

You can use the `chmod` command to set permissions in either of two modes:

- **Absolute Mode** – Use numbers to represent file permissions (the method most commonly used to set permissions). When you change permissions by using the absolute mode, you represent permissions for each triplet by an octal mode number.
- **Symbolic Mode** – Use combinations of letters and symbols to add or remove permissions.

The following table lists the octal values for setting file permissions in absolute mode. You use these numbers in sets of three to set permissions for owner, group, and other (in that order). For example, the value 644 sets read and write permissions for owner, and read-only permissions for group and other.

**TABLE 15-5** Setting File Permissions in Absolute Mode

Octal Value	File Permissions Set	Permissions Description
0	---	No permissions
1	--x	Execute permission only
2	-w-	Write permission only
3	-wx	Write and execute permissions
4	r--	Read permission only
5	r-x	Read and execute permissions
6	rw-	Read and write permissions
7	rwx	Read, write, and execute permissions

You can set special permissions on a file in absolute or symbolic modes. However, you cannot set or remove `setuid` permissions on a directory by using absolute mode. You must use symbolic mode. In absolute mode, you set special permissions by adding a new octal value to the left of the permission triplet. The following table lists the octal values to set special permissions on a file.

**TABLE 15-6** Setting Special Permissions in Absolute Mode

Octal Value	Special Permissions Set
1	Sticky bit
2	setgid
4	setuid

The following table lists the symbols for setting file permissions in symbolic mode. Symbols can specify whose permissions are to be set or changed, the operation to be performed, and the permissions that are being assigned or changed.

**TABLE 15-7** Setting File Permissions in Symbolic Mode

Symbol	Function	Description
u	Who	User (owner)
g	Who	Group
o	Who	Others
a	Who	All
=	Operator	Assign
+	Operator	Add
-	Operator	Remove
r	Permission	Read
w	Permission	Write
x	Permission	Execute
l	Permission	Mandatory locking, setgid bit is on, group execution bit is off
s	Permission	setuid or setgid bit is on
S	Permission	suid bit is on, user execution bit is off
t	Permission	Sticky bit is on, execution bit for others is on
T	Permission	Sticky bit is on, execution bit for others is off

The *who operator permission* designations in the function column specifies the symbols that change the permissions on the file or directory.

*who* Specifies whose permissions are to be changed.

<i>operator</i>	Specifies the operation to be performed.
<i>permissions</i>	Specifies what permissions are to be changed.

## ▼ How to Change Permissions in Absolute Mode

Use the following procedure to change permissions in absolute mode.

1. **If you are not the owner of the file or directory, become superuser or assume an equivalent role.**

Only the current owner or superuser can use the `chmod` command to change file permissions on a file or directory.

2. **Change permissions in absolute mode by using the `chmod` command.**

```
$ chmod nnn filename
```

*nnn* Specifies the octal values that represent the permissions for the file owner, file group, and others, in that order. See Table 15-5 for the list of valid octal values.

*filename* Specifies the file or directory.

---

**Note** – If you use the `chmod` command to change the file group permissions on a file with ACL entries, both the file group permissions and the ACL mask are changed to the new permissions. Be aware that the new ACL mask permissions can change the permissions for additional users and groups who have ACL entries on the file. Use the `getfacl` command to make sure that the appropriate permissions are set for all ACL entries. For more information, see the `getfacl(1)` man page.

---

3. **Verify that the permissions of the file have changed.**

```
$ ls -l filename
```

### Example—Changing Permissions in Absolute Mode

In the following example, the permissions of a public directory are changed from 744 (read, write, execute; read-only; and read-only) to 755 (read, write, execute; read and execute; and read and execute).

```
$ ls -ld public_dir
drwxr--r-- 1 ignatz  staff    6023 Aug  5 12:06 public_dir
$ chmod 755 public_dir
```

```
$ ls -ld public_dir
drwxr-xr-x 1 ignatz  staff    6023 Aug  5 12:06 public_dir
```

In the following example, the permissions of an executable shell script are changed from read and write to read, write, and execute.

```
$ ls -l my_script
-rw----- 1 ignatz  staff    6023 Aug  5 12:06 my_script
$ chmod 700 my_script
$ ls -l my_script
-rwx----- 1 ignatz  staff    6023 Aug  5 12:06 my_script
```

## ▼ How to Change Special Permissions in Absolute Mode

Use the following procedure to change special permissions in absolute mode.

1. **If you are not the owner of the file or directory, become superuser or assume an equivalent role.**

Only the current owner or superuser can use the `chmod` command to change the special permissions on a file or directory.

2. **Change special permissions in absolute mode by using the `chmod` command.**

```
$ chmod nnnn filename
```

<i>nnnn</i>	Specifies the octal values that change the permissions on the file or directory. The first octal value on the left sets the special permissions on the file. For the list of valid octal values for the special permissions, see Table 15-6.
-------------	--

<i>filename</i>	Specifies the file or directory.
-----------------	----------------------------------

---

**Note** – If you use the `chmod` command to change the file group permissions on a file with ACL entries, both the file group permissions and the ACL mask are changed to the new permissions. Be aware that the new ACL mask permissions can change the permissions for additional users and groups who have ACL entries on the file. Use the `getfacl` command to make sure that the appropriate permissions are set for all ACL entries. For more information, see the `getfacl(1)` man page.

---

3. **Verify that the permissions of the file have changed.**

```
$ ls -l filename
```

## Examples—Setting Special Permissions in Absolute Mode

In the following example, the `setuid` permission is set on the `dbprog` file.

```
$ chmod 4555 dbprog
$ ls -l dbprog
-r-sr-xr-x  1 db      staff      12095 May  6 09:29 dbprog
```

In the following example, the `setgid` permission is set on the `dbprog2` file.

```
$ chmod 2551 dbprog2
$ ls -l dbprog2
-r-xr-s--x  1 db      staff      24576 May  6 09:30 dbprog2
```

In the following example, the sticky bit permission is set on the `public_dir` directory.

```
$ chmod 1777 public_dir
$ ls -ld public_dir
drwxrwxrwt  2 ignatz  staff      512 May 15 15:27 public_dir
```

## ▼ How to Change Permissions in Symbolic Mode

Use the following procedure to change permissions in symbolic mode.

- 1. If you are not the owner of the file or directory, become superuser.**  
Only the current owner or superuser can use the `chmod` command to change file permissions on a file or directory.
- 2. Change permissions in symbolic mode by using the `chmod` command.**

```
$ chmod who operator permission filename
```

*who operator permission*      *who* specifies whose permissions are to be changed, *operator* specifies the operation to be performed, and *permission* specifies what permissions are to be changed. For the list of valid symbols, see Table 15-7.

*filename*                      Specifies the file or directory.

- 3. Verify that the permissions of the file have changed.**

```
$ ls -l filename
```

## Examples—Changing Permissions in Symbolic Mode

In the following example, read permission are taken away from others.

```
$ chmod o-r filea
```

In the following example, read and execute permissions are added for user, group, and others.

```
$ chmod a+rx fileb
```

In the following example, read, write, and execute permissions are assigned to group.

```
$ chmod g=rwx filec
```

---

## Searching for Special Permissions

You should monitor your system for any unauthorized use of the `setuid` and `setgid` permissions on programs to gain superuser privileges. A suspicious listing grants ownership of such a program to a user rather than to `root` or `bin`.

### ▼ How to Find Files With `setuid` Permissions

Use the following procedure to find files with `setuid` permissions.

1. **Become superuser or assume an equivalent role.**
2. **Find files with `setuid` permissions by using the `find` command.**

```
# find directory -user root -perm -4000 -exec ls -ldb {} \; >/tmp/filename
```

<code>find <i>directory</i></code>	Checks all mounted paths starting at the specified <i>directory</i> , which can be <code>root (/)</code> , <code>sys</code> , <code>bin</code> , or <code>mail</code> .
<code>-user root</code>	Displays files owned only by <code>root</code> .
<code>-perm -4000</code>	Displays files only with permissions set to 4000.
<code>-exec ls -ldb</code>	Displays the output of the <code>find</code> command in <code>ls -ldb</code> format.
<code>&gt;/tmp/<i>filename</i></code>	Writes results to this file.

3. **Display the results in `/tmp/filename`.**

```
# more /tmp/filename
```

If you need background information about `setuid` permissions, see “`setuid` Permission” on page 213.



## Example—Finding Files With `setuid` Permissions

```
# find / -user root -perm -4000 -exec ls -ldb {} \; > /tmp/ckprm
# cat /tmp/ckprm
-r-sr-xr-x 1 root bin 38836 Aug 10 16:16 /usr/bin/at
-r-sr-xr-x 1 root bin 19812 Aug 10 16:16 /usr/bin/crontab
---s--x--x 1 root sys 46040 Aug 10 15:18 /usr/bin/ct
-r-sr-xr-x 1 root sys 12092 Aug 11 01:29 /usr/lib/mv_dir
-r-sr-sr-x 1 root bin 33208 Aug 10 15:55 /usr/lib/lpadmin
-r-sr-sr-x 1 root bin 38696 Aug 10 15:55 /usr/lib/lpsched
---s--x--- 1 root rar 45376 Aug 18 15:11 /usr/rar/bin/sh
-r-sr-xr-x 1 root bin 12524 Aug 11 01:27 /usr/bin/df
-rwsr-xr-x 1 root sys 21780 Aug 11 01:27 /usr/bin/newgrp
-r-sr-sr-x 1 root sys 23000 Aug 11 01:27 /usr/bin/passwd
-r-sr-xr-x 1 root sys 23824 Aug 11 01:27 /usr/bin/su
```

This output shows that a user named `rar` has made a personal copy of `/usr/bin/sh`, and has set the permissions as `setuid` to root. As a result, `rar` can execute `/usr/rar/bin/sh` and become the privileged user. If you want to save this output for future reference, move the file out of the `/tmp` directory.

---

## Executable Stacks and Security

A number of security bugs are related to default executable stacks when their permissions are set to read, write, and execute. While stacks with execute permissions are allowed, most programs can function correctly without using executable stacks.

The `noexec_user_stack` variable (available starting in the Solaris 2.6 release) enables you to specify whether stack mappings are executable. By default, the variable is set to zero, except on 64-bit applications, which provides ABI-compliant behavior. If the variable is set to non-zero, the system marks the stack of every process in the system as readable and writable, but not executable.

Once this variable is set, programs that attempt to execute code on their stack are sent a `SIGSEGV` signal, which usually results in the program terminating with a core dump. Such programs also generate a warning message that includes the name of the offending program, the process ID, and the real UID of the user who ran the program. For example:

```
a.out[347] attempt to execute code on stack by uid 555
```

The message is logged by the `syslog` daemon when the `syslog` kern facility is set to `notice` level. This logging is set by default in the `syslog.conf` file, which means that the message is sent to both the console and to the `/var/adm/messages` file. For more information, see the `syslogd(1M)` and `syslog.conf(4)` man pages.

This message is useful for observing potential security problems, as well as to identify valid programs that depend upon executable stacks that have been prevented from correct operation by setting this variable. If the administrator does not want any messages logged, then the `noexec_user_stack_log` variable can be set to zero to disable it in the `/etc/system` file, though the `SIGSEGV` signal can continue to cause the executing program to core dump.

You can use `mprotect` if you want programs to explicitly mark their stack as executable. For more information, see the `mprotect(2)` man page.

Because of hardware limitations, the capability of catching and reporting executable stack problems is only available on `sun4m` and `sun4u` platforms.

## ▼ How to Disable Programs From Using Executable Stacks

1. **Become superuser or assume an equivalent role.**
2. **Edit the `/etc/system` file and add the following line:**

```
set noexec_user_stack=1
```

3. **Reboot the system.**

```
# init 6
```

## ▼ How to Disable Executable Stack Message Logging

1. **Become superuser or assume an equivalent role.**
2. **Edit the `/etc/system` file and add the following line:**

```
set noexec_user_stack_log=0
```

3. **Reboot the system.**

```
# init 6
```

---

## Using Access Control Lists (ACLs)

Traditional UNIX file protection provides read, write, and execute permissions for the three user classes: file owner, file group, and other. An ACL provides better file security by enabling you to define file permissions for the file owner, file group, other, specific users and groups, and default permissions for each of those categories.

For example, if you wanted everyone in a group to be able to read a file, you would simply give group read permissions on that file. Now, assume you wanted only one person in the group to be able to write to that file. Standard UNIX doesn't provide that level of file security. However, this dilemma is perfect for ACLs.

ACL entries are the way to define an ACL on a file, and they are set through the `setfacl` command. ACL entries consist of the following fields separated by colons:

*entry-type*: *[uid | gid]* :*perms*

<i>entry-type</i>	Is the type of ACL entry on which to set file permissions. For example, <i>entry-type</i> can be <code>user</code> (the owner of a file) or <code>mask</code> (the ACL mask). For a listing of ACL entries, see Table 15-8 and Table 15-9.
<i>uid</i>	Is the user name or user ID (UID).
<i>gid</i>	Is the group name or group ID (GID).
<i>perms</i>	Represents the permissions that are set on <i>entry-type</i> . <i>perms</i> can be indicated by the symbolic characters <code>rxw</code> or a number (the same permissions numbers that are used with the <code>chmod</code> command).

The following example shows an ACL entry that sets read and write permissions for the user `nathan`.

```
user:nathan:rw-
```



---

**Caution** – UFS file system attributes such as ACLs are supported in UFS file systems only. Thus, if you restore or copy files with ACL entries into the `/tmp` directory, which is usually mounted as a TMPFS file system, the ACL entries will be lost. Use the `/var/tmp` directory for temporary storage of UFS files.

---

## ACL Entries for Files

The following table lists the valid ACL entries that you might use when setting file ACLs. The first three ACL entries provide the basic UNIX file protection.

**TABLE 15-8** ACL Entries for Files

ACL Entry	Description
<code>u[ser]::perms</code>	File owner permissions.
<code>g[roup]::perms</code>	File group permissions.

**TABLE 15–8** ACL Entries for Files (Continued)

ACL Entry	Description
<code>o[ther] :perms</code>	Permissions for users other than the file owner or members of the file group.
<code>m[ask] :perms</code>	The ACL mask. The mask entry indicates the maximum permissions that are allowed for users (other than the owner) and for groups. The mask is a quick way to change permissions on all the users and groups.  For example, the <code>mask : r - -</code> mask entry indicates that users and groups cannot have more than read permissions, even though they might have write and execute permissions.
<code>u[ser] :uid:perms</code>	Permissions for a specific user. For <i>uid</i> , you can specify either a user name or a numeric UID.
<code>g[roup] :gid:perms</code>	Permissions for a specific group. For <i>gid</i> , you can specify either a group name or a numeric GID.

## ACL Entries for Directories

In addition to the ACL entries that are described in Table 15–8, you can set default ACL entries on a directory. Files or directories created in a directory that has default ACL entries will have the same ACL entries as the default ACL entries. Table 15–9 lists the default ACL entries for directories.

When you set default ACL entries for specific users and groups on a directory for the first time, you must also set default ACL entries for the file owner, file group, others, and the ACL mask. These entries are required and are the first four default ACL entries in the following table.

**TABLE 15–9** Default ACL Entries for Directories

Default ACL Entry	Description
<code>d[efault] :u[ser] ::perms</code>	Default file owner permissions.
<code>d[efault] :g[roup] ::perms</code>	Default file group permissions.
<code>d[efault] :o[ther] :perms</code>	Default permissions for users other than the file owner or members of the file group.
<code>d[efault] :m[ask] :perms</code>	Default ACL mask.

**TABLE 15–9** Default ACL Entries for Directories (Continued)

Default ACL Entry	Description
<code>d[efault]:u[ser]:uid:perms</code>	Default permissions for a specific user. For <i>uid</i> , you can specify either a user name or a numeric UID.
<code>d[efault]:g[roup]:gid:perms</code>	Default permissions for a specific group. For <i>gid</i> , you can specify either a group name or a numeric GID.

## ▼ How to Set an ACL on a File

### 1. Set an ACL on a file by using the `setfacl` command.

```
$ setfacl -s user::perms,group::perms,other:perms,mask:perms,acl-entry-list filename ...
```

<code>-s</code>	Sets an ACL on the file. If a file already has an ACL, it is replaced. This option requires at least the file owner, file group, and other entries.
<code>user::perms</code>	Specifies the file owner permissions.
<code>group::perms</code>	Specifies the file group permissions.
<code>other:perms</code>	Specifies the permissions for users other than the file owner or members of the file group.
<code>mask:perms</code>	Specifies the permissions for the ACL mask. The mask indicates the maximum permissions that are allowed for users (other than the owner) and for groups.
<code>acl-entry-list</code>	Specifies the list of one or more ACL entries to set for specific users and groups on the file or directory. You can also set default ACL entries on a directory. Table 15–8 and Table 15–9 show the valid ACL entries.
<code>filename ...</code>	Specifies one or more files or directories on which to set the ACL, separated by a space.



---

**Caution** – If an ACL already exists on the file, the `-s` option will replace the entire ACL with the new ACL.

---

For more information, see `setfacl(1)`.

### 2. Verify that an ACL was set on the file or verify which ACL entries were set on the file.

```
$ getfacl filename
```

For more information, see “How to Check If a File Has an ACL” on page 231.

## Examples—Setting an ACL on a File

In the following example, the file owner permissions are set to read and write, file group permissions are set to read only, and other permissions are set to none on the `ch1.doc` file. In addition, the user `george` is given read and write permissions on the file, and the ACL mask permissions are set to read and write, which means that no user or group can have execute permissions.

```
$ setfacl -s user::rw-,group::r--,other:---,mask:rw-,user:george:rw- ch1.doc
$ ls -l
total 124
-rw-r-----+ 1 nathan  sysadmin   34816 Nov 11 14:16 ch1.doc
-rw-r--r--   1 nathan  sysadmin   20167 Nov 11 14:16 ch2.doc
-rw-r--r--   1 nathan  sysadmin    8192 Nov 11 14:16 notes
$ getfacl ch1.doc
# file: ch1.doc
# owner: nathan
# group: sysadmin
user::rw-
user:george:rw-   #effective:rw-
group::r--        #effective:r--
mask:rw-
other:---
```

In the following example, the file owner permissions are set to read, write, and execute, file group permissions are set to read only, other permissions are set to none, and the ACL mask permissions are set to read on the `ch2.doc` file. In addition, the user `george` is given read and write permissions. However, due to the ACL mask, the permissions for `george` are read only.

```
$ setfacl -s u::7,g::4,o::0,m:4,u:george:7 ch2.doc
$ getfacl ch2.doc
# file: ch2.doc
# owner: nathan
# group: sysadmin
user::rwx
user:george:rwx   #effective:r--
group::r--        #effective:r--
mask:r--
other:---
```

## ▼ How to Copy an ACL

Copy a file's ACL to another file by redirecting the `getfacl` output.

```
$ getfacl filename1 | setfacl -f - filename2
```

*filename1* Specifies the file from which to copy the ACL.

*filename2* Specifies the file on which to set the copied ACL.

## Example—Copying an ACL

In the following example, the ACL on `ch2.doc` is copied to `ch3.doc`.

```
$ getfacl ch2.doc | setfacl -f - ch3.doc
```

## ▼ How to Check If a File Has an ACL

Check if a file has an ACL by using the `ls` command.

```
$ ls -l filename
```

*filename* specifies the file or directory.

In the output, a plus sign (+) to the right of the mode field indicates that the file has an ACL.

---

**Note** – Unless you have added ACL entries for additional users or groups on a file, a file is considered to be a “trivial” ACL and the plus sign (+) will not display.

---

## Example—Checking If a File Has an ACL

The following example shows that the `ch1.doc` file has an ACL, because the listing has a plus sign (+) to the right of the mode field.

```
$ ls -l ch1.doc
-rwxr-----+ 1 nathan  sysadmin    167 Nov 11 11:13 ch1.doc
```

## ▼ How to Modify ACL Entries on a File

### 1. Modify ACL entries on a file by using the `setfacl` command.

```
$ setfacl -m acl-entry-list filename ...
```

<code>-m</code>	Modifies the existing ACL entry.
<code>acl-entry-list</code>	Specifies the list of one or more ACL entries to modify on the file or directory. You can also modify default ACL entries on a directory. Table 15–8 and Table 15–9 show the valid ACL entries.
<code>filename ...</code>	Specifies one or more files or directories, separated by a space.

### 2. Verify that the ACL entries were modified on the file by using the `getfacl` command.

```
$ getfacl filename
```

## Examples—Modifying ACL Entries on a File

In the following example, the permissions for the user `george` are modified to read and write.

```
$ setfacl -m user:george:6 ch3.doc
$ getfacl ch3.doc
# file: ch3.doc
# owner: nathan
# group: staff
user::rw-
user::george:rw-          #effective:r--
group::r-                 #effective:r--
mask:r--
other:r-
```

In the following example, the default permissions for the group `staff` are modified to read and the default ACL mask permissions are modified to read and write on the `book` directory.

```
$ setfacl -m default:group:staff:4,default:mask:6 book
```



## ▼ How to Delete ACL Entries From a File

1. Delete ACL entries from a file by using the `setfacl` command.

```
$ setfacl -d acl-entry-list filename ...
```

<code>-d</code>	Deletes the specified ACL entries.
<code>acl-entry-list</code>	Specifies the list of ACL entries (without specifying the permissions) to delete from the file or directory. You can only delete ACL entries and default ACL entries for specific users and groups. Table 15-8 and Table 15-9 show the valid ACL entries.
<code>filename ...</code>	Specifies one or more files or directories, separated by a space.

Alternately, you can use `setfacl -s` to delete all the ACL entries on a file and replace them with the new ACL entries that are specified.

2. To verify that the ACL entries were deleted from the file, by using the `getfacl` command.

```
$ getfacl filename
```

### Example—Deleting ACL Entries on a File

In the following example, the user `george` is deleted from the `ch4.doc` file.

```
$ setfacl -d user:george ch4.doc
```

## ▼ How to Display ACL Entries for a File

Display ACL entries for a file by using the `getfacl` command.

```
$ getfacl [-a | -d] filename ...
```

<code>-a</code>	Displays the file name, file owner, file group, and ACL entries for the specified file or directory.
<code>-d</code>	Displays the file name, file owner, file group, and default ACL entries for the specified directory.
<code>filename ...</code>	Specifies one or more files or directories, separated by a space.

If you specify multiple file names on the command line, the ACL entries are displayed with a blank line between each entry.

## Examples—Displaying ACL Entries for a File

The following example shows all the ACL entries for the `ch1.doc` file. The `#effective:` note beside the user and group entries indicates what the permissions are after being modified by the ACL mask.

```
$ getfacl ch1.doc

# file: ch1.doc
# owner: nathan
# group: sysadmin
user::rw-
user:george:r--      #effective:r--
group::rw-          #effective:rw-
mask:rw-
other:---
```

The following example shows the default ACL entries for the `book` directory.

```
$ getfacl -d book

# file: book
# owner: nathan
# group: sysadmin
user::rwx
user:george:r-x      #effective:r-x
group::rwx           #effective:rwx
mask:rwx
other:---
default:user::rw-
default:user:george:r--
default:group::rw-
default:mask:rw-
default:other:---
```

---

## Securing Systems (Tasks)

---

This chapter describes the procedures for securing systems in the Solaris environment.

This is a list of the step-by-step instructions in this chapter.

- “How to Display a User’s Login Status” on page 235
- “How to Display Users Without Passwords” on page 236
- “How to Temporarily Disable User Logins” on page 237
- “How to Save Failed Login Attempts” on page 238
- “How to Create a Dial-up Password” on page 240
- “How to Temporarily Disable Dial-up Logins” on page 241
- “How to Restrict Superuser (root) Login to the Console” on page 242
- “How to Monitor Who Is Using the su Command” on page 243
- “How to Display Superuser (root) Access Attempts to the Console” on page 243
- “How to Disable or Enable a System’s Abort Sequence” on page 244

For overview information about securing systems, see “System Security” on page 202.

---

## Displaying User Login Information

This section describes how to display user login information.

### ▼ How to Display a User’s Login Status

1. **Become superuser or assume an equivalent role.**
2. **Display a user’s login status by using the `logins` command.**

```
# logins -x -l username
```

- x Displays an extended set of login status information.
- l *username* Displays the login status for the specified user. *username* is a user's login name. Multiple login names must be specified in a comma-separated list.

The `logins` command uses the local `/etc/passwd` file and the name service password databases (if any) to obtain a user's login status. For more information, see the `logins(1M)` man page.

## Example—Displaying a User's Login Status

In the following example, the login status for the user `rimmer` is displayed.

```
# logins -x -l rimmer
rimmer      500      staff          10   Arnold J. Rimmer
            /export/home/rimmer
            /bin/sh
            PS 010170 10 7 -1
```

rimmer	Identifies the user's login name.
500	Identifies the user ID (UID).
staff	Identifies the user's primary group.
10	Identifies the group ID (GID).
Arnold J. Rimmer	Identifies the comment.
/export/home/rimmer	Identifies the user's home directory.
/bin/sh	Identifies the login shell.
PS 010170 10 7 -1	Specifies the password aging information: <ul style="list-style-type: none"> <li>■ Last date the password was changed</li> <li>■ Number of days that are required between changes</li> <li>■ Number of days that are allowed before a change is required</li> <li>■ Warning period</li> </ul>

## ▼ How to Display Users Without Passwords

You should make sure that all users have a valid password.

1. **Become superuser or assume an equivalent role.**
2. **Display all users who have no passwords by using the `logins` command.**

```
# logins -p
```

The `-p` option displays a list of users with no passwords. The `logins` command uses the local `/etc/passwd` file and the name service password databases to obtain a user's login status.

## Example—Displaying Users Without Passwords

The following example shows that the user `pmorph` does not have a password.

```
# logins -p
pmorph      501      other          1      Polly Morph
#
```

---

## Temporarily Disabling User Logins

You can temporarily disable user logins by doing either of the following:

- Creating the `/etc/nologin` file.
- Bringing the system to run level 0 (single-user mode). For information on bringing the system to single-user mode, see “Shutting Down a System (Tasks)” in *System Administration Guide: Basic Administration*.

### Creating the `/etc/nologin` File

Create this file to disallow user logins and notify users when a system will be unavailable for an extended period of time because of a system shutdown or routine maintenance.

If a user attempts to log in to a system where this file exists, the contents of the `nologin` file is displayed, and the user login is terminated. Superuser logins are not affected. For more information, see the `nologin(4)` man page.

## ▼ How to Temporarily Disable User Logins

1. **Become superuser or assume an equivalent role.**
2. **Create the `/etc/nologin` file by using an editor.**

```
# vi /etc/nologin
```

3. **Include a message about system availability.**

4. Close and save the file.

## Example—Disabling User Logins

This example shows how to notify users of system unavailability.

```
# vi /etc/nologin
(Add system message here)

# cat /etc/nologin
***No logins permitted.***

***The system will be unavailable until 12 noon.***
```

---

## Saving Failed Login Attempts

You can save failed login attempts by creating the `/var/adm/loginlog` file with read and write permission for `root` only. After you create the `loginlog` file, all failed login activity will be automatically written to this file after five failed attempts. For detailed instructions, see “How to Save Failed Login Attempts” on page 238.

The `loginlog` file contains one entry for each failed attempt. Each entry contains the user’s login name, tty device, and time of the failed attempt. If a person makes fewer than five unsuccessful attempts, none of the attempts are logged.

The `loginlog` file might grow quickly. To use the information in this file and to prevent the file from getting too large, you must check and clear its contents occasionally. If this file shows a lot of activity, it might suggest an attempt to break into the computer system. For more information about this file, see the `loginlog(4)` man page.

### ▼ How to Save Failed Login Attempts

1. Become superuser or assume an equivalent role.
2. Create the `loginlog` file in the `/var/adm` directory.

```
# touch /var/adm/loginlog
```
3. Set read and write permissions for `root` on the `loginlog` file.

```
# chmod 600 /var/adm/loginlog
```
4. Change group membership to `sys` on the `loginlog` file.

```
# chgrp sys /var/adm/loginlog
```

5. Make sure that the log works by attempting to log into the system five times with the wrong password. Then, display the `/var/adm/loginlog` file.

```
# more /var/adm/loginlog
rimmer:/dev/pts/1:Wed Jan 16 09:22:31 2002
rimmer:/dev/pts/1:Wed Jan 16 09:22:39 2002
rimmer:/dev/pts/1:Wed Jan 16 09:22:45 2002
rimmer:/dev/pts/1:Wed Jan 16 09:22:53 2002
rimmer:/dev/pts/1:Wed Jan 16 09:23:01 2002
#
```

---

## Password Protection Using Dial-up Passwords

You can add a layer of security to your password protection by requiring a *dial-up password* for users who access a system through a modem or dial-up port. A dial-up password is an additional password that a user must enter before being granted access to the system.

Only superuser can create or change a dial-up password. To ensure the integrity of the system, the password should be changed about once a month. The most effective use of this feature is to require a dial-up password to gain access to a gateway system.

Two files are involved in creating a dial-up password, `/etc/dialups` and `/etc/d_passwd`. The first file contains a list of ports that require a dial-up password. The second file contains a list of shell programs that require an encrypted password as the additional dial-up password.

If the user's login program (as specified in `/etc/passwd`) is not found in `/etc/d_passwd`, or if the login shell field in `/etc/passwd` is null, the password entry for `/usr/bin/sh` is used.

- If the user's login shell in `/etc/passwd` matches an entry in `/etc/d_passwd`, the user must supply a dial-up password.
- If the user's login shell in `/etc/passwd` is not found in `/etc/d_passwd`, the user must supply the default password. The default password is the entry for `/usr/bin/sh`.
- If the login shell field in `/etc/passwd` is empty, the user must supply the default password (the entry for `/usr/bin/sh`).
- If `/etc/d_passwd` has no entry for `/usr/bin/sh`, then those users whose login shell field in `/etc/passwd` is empty or does not match any entry in `/etc/d_passwd` will not be prompted for a dial-up password.

- Dial-up logins are disabled if `/etc/d_passwd` has only the following entry:  
`/usr/bin/sh:*:`

## ▼ How to Create a Dial-up Password



---

**Caution** – When you first establish a dial-up password, be sure to remain logged in on at least one port while you are testing the password on a different port. If you make a mistake while you are installing the extra password and log off to test the new password, you might not be able to log back on. If you are still logged in on another port, you can go back and fix your mistake.

---

1. **Become superuser or assume an equivalent role.**
2. **Create an `/etc/dialups` file that contains a list of serial devices, including all the ports that will require dial-up password protection.**

The `/etc/dialups` file should look like the following:

```
/dev/term/a
/dev/term/b
/dev/term/c
```

3. **Create an `/etc/d_passwd` file that contains the login programs that will require a dial-up password, and the encrypted dial-up password.**

Include shell programs that a user could be running at login, for example, `uucico`, `sh`, `ksh`, and `csh`. The `/etc/d_passwd` file should look like the following:

```
/usr/lib/uucp/uucico:encrypted-password:
/usr/bin/csh:encrypted-password:
/usr/bin/ksh:encrypted-password:
/usr/bin/sh:encrypted-password:
```

4. **Set ownership to `root` on the two files.**  

```
# chown root /etc/dialups /etc/d_passwd
```
5. **Set group ownership to `root` on the two files.**  

```
# chgrp root /etc/dialups /etc/d_passwd
```
6. **Set read and write permissions for `root` on the two files.**  

```
# chmod 600 /etc/dialups /etc/d_passwd
```
7. **Create the encrypted passwords.**

- a. **Create a temporary user.**

```
# useradd username
```



**b. Create a password for the temporary user.**

```
# passwd username
```

**c. Capture the encrypted password.**

```
# grep username /etc/shadow > username.temp
```

**d. Edit the `username.temp` file.**

Delete all fields except the encrypted password (the second field).

For example, in the following line, the encrypted password is `U9gp9SyA/J1Sk`.

```
temp:U9gp9SyA/J1Sk:7967::::::::7988:
```

**e. Delete the temporary user.**

```
# userdel username
```

**8. Copy the encrypted password from `username.temp` file into the `/etc/d_passwd` file.**

You can create a different password for each login shell, or use the same password for each login shell.

## ▼ How to Temporarily Disable Dial-up Logins

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

**2. Put the following entry by itself into the `/etc/d_passwd` file:**

```
/usr/bin/sh:*:
```

---

## Restricting Superuser (root) Access on the Console

The superuser account is used by the operating system to accomplish basic functions, and has wide-ranging control over the entire operating system. The superuser account has access to and can execute essential system programs. For this reason, there are almost no security restraints for any program that is run by superuser.

You can protect the superuser account on a system by restricting access to a specific device through the `/etc/default/login` file. For example, if superuser access is restricted to the console, you can log in to a system as superuser only from the console. If anybody remotely logs in to the system to perform an administrative function, they must first log in with their user login and then use the `su` command to become superuser. See the following section for detailed instructions.

---

**Note** – Restricting superuser login to the console is set up by default when you install the Solaris release.

---

An alternative to using the superuser account is to setup role-based access control (RBAC). For overview information on RBAC, see Chapter 17.

## ▼ How to Restrict Superuser (root) Login to the Console

1. **Become superuser or assume an equivalent role.**
2. **Edit the `/etc/default/login` file.**
3. **Uncomment the following line:**

```
CONSOLE=/dev/console
```

Any users who try to remotely log in to this system must first log in with their user login, and then use the `su` command to become superuser.

4. **Attempt to log in remotely as superuser to this system, and verify that the operation fails.**

---

## Monitoring Who Is Using the `su` Command

You can monitor `su` attempts by monitoring the `/var/adm/sulog` file. This file logs each time the `su` command is used. The `su` logging in this file is enabled by default through the following entry in the `/etc/default/su` file:

```
SULOG=/var/adm/sulog
```

The `sulog` file lists all uses of the `su` command, not only the `su` attempts that are used to switch from user to superuser. The entries show the date and time the command was entered, whether or not the attempt was successful (+ or -), the port from which the command was issued, and finally, the name of the user and the switched identity.

Through the `/etc/default/su` file, you can set up the system to display on the console each time an attempt is made to use the `su` command to gain superuser access from a remote system. This method is a good way to immediately detect someone who is trying to gain superuser access on the system that you are currently working on. See the following section for detailed instructions.

## ▼ How to Monitor Who Is Using the `su` Command

1. **Become superuser or assume an equivalent role.**
2. **Monitor the contents of the `/var/adm/sulog` file on a regular basis.**

```
# more /var/adm/sulog
SU 12/20 16:26 + pts/0 nathan-root
SU 12/21 10:59 + pts/0 nathan-root
SU 01/12 11:11 + pts/0 root-joebob
SU 01/12 14:56 + pts/0 pmorph-root
SU 01/12 14:57 + pts/0 pmorph-root
```

## ▼ How to Display Superuser (`root`) Access Attempts to the Console

1. **Become superuser or assume an equivalent role.**
2. **Edit the `/etc/default/su` file.**
3. **Uncomment the following line:**

```
CONSOLE=/dev/console
```

Use the `su` command to become `root`, and verify that a message is printed on the system console.

---

## Modifying a System's Abort Sequence

Use the following procedure to disable or enable a system's abort sequence. The default system behavior is that a system's abort sequence is enabled.

Some server systems have a key switch that if set in the secure position, overrides the software keyboard abort settings. So, any changes you make with the following procedure might not be implemented.

## ▼ How to Disable or Enable a System's Abort Sequence

1. Become superuser or assume an equivalent role.
2. Select one of the following to disable or enable a system's abort sequence:
  - a. Remove the pound sign (#) from the following line in the `/etc/default/kbd` file to disable a system's abort sequence:

```
#KEYBOARD_ABORT=disable
```

- b. Add the pound sign (#) to the following line in the `/etc/default/kbd` file to enable a system's abort sequence:

```
KEYBOARD_ABORT=disable
```

3. Update the keyboard defaults.

```
# kbd -i
```

## Role-Based Access Control (Overview)

---

This chapter describes role-based access control (RBAC), a security feature for controlling access to tasks that would normally be restricted to superuser. This is a list of the overview information in this chapter.

- “RBAC: Replacing the Superuser Model” on page 245
- “Solaris RBAC Elements” on page 246
- “Privileged Applications” on page 248
- “RBAC Roles” on page 250
- “RBAC Authorizations” on page 250
- “RBAC Rights Profiles” on page 251
- “Name Service Scope” on page 251

For information on RBAC tasks, see Chapter 18. For reference information on the RBAC elements and tools, see Chapter 19.

---

### RBAC: Replacing the Superuser Model

In conventional UNIX systems, the `root` user (also referred to as superuser) is all-powerful, with the ability to read and write to any file, run all programs, and send kill signals to any process. Effectively, anyone who can become superuser can modify a site’s firewall, alter the audit trail, read payroll and other confidential records, and shut down the entire network.

Role-based access control (RBAC) is an alternative to the all-or-nothing superuser model. RBAC uses the security principle of least privilege, which is that no user should be given more privilege than necessary for performing his or her job. RBAC allows an organization to separate superuser’s capabilities and assign them to special user accounts that are called *roles*. Roles can be assigned to specific individuals, according to their job needs.

The flexibility in setting up roles enables a variety of security policies. Three recommended roles that can be easily configured are available:

- **Primary Administrator** – A powerful role that is equivalent to `root`.
- **System Administrator** – A less strong role for administration that is not related to security. This role does not allow the user to set passwords.
- **Operator** – A junior administrator role for operations such as backups and restores, and printer management.

There is no requirement that these specific roles be implemented. Roles are a function of an organization's security needs. Roles can be set up for special-purpose administrators in areas such as security, networking, or firewall administration. Another strategy is to create a single strong administrator role along with an advanced user role for those users who are permitted to fix portions of their own systems.


---

## Solaris RBAC Elements

In the RBAC model in the Solaris operating environment, users log in as themselves and assume roles that enable them to run restricted administration graphical tools and commands. The RBAC model introduces these elements to the Solaris operating environment:

- **Privileged Application** – An application that can override system controls and check for specific user IDs (UIDs), group IDs (GIDs), or authorizations (see “Privileged Applications” on page 248).
- **Role** – A special identity for running privileged applications that can be assumed by assigned users only.
- **Authorization** – A permission that can be assigned to a role or user (or embedded in a rights profile) for performing a class of actions that are otherwise prohibited by security policy.
- **Rights Profile** – A collection of overrides that can be assigned to a role or user. A rights profile can consist of authorizations, commands with `setuid` or `setgid` permissions (referred to as security attributes), and other rights profiles.

The following figure shows how the RBAC elements work together.



**FIGURE 17-1** Solaris RBAC Element Relationships

In RBAC, users are assigned to roles. Roles get their capabilities from rights profiles and authorizations. Authorizations are generally assigned to the rights profiles with which they are logically associated but can be assigned directly to roles.

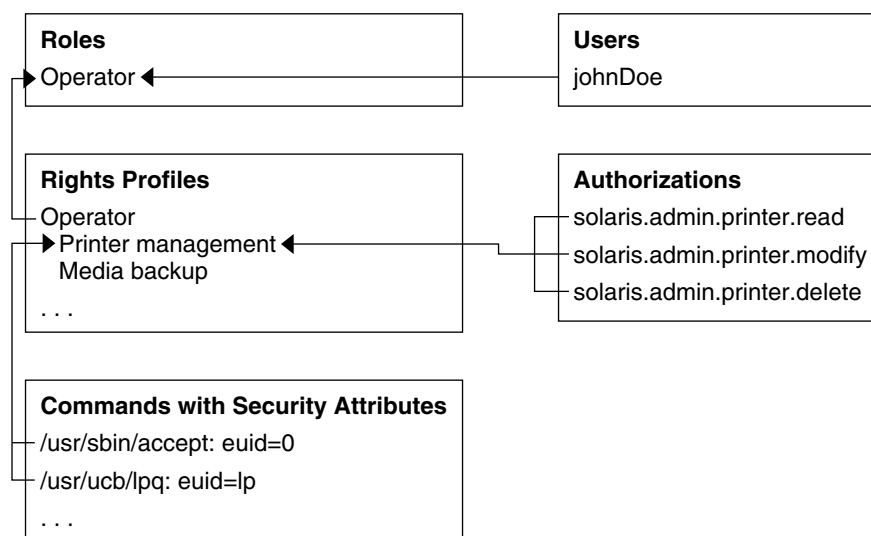
---

**Note** – Rights profiles and authorizations can also be assigned directly to users. This practice is discouraged because it enables users to make mistakes through inadvertent use of their privileges.

---

Commands with security attributes, that is, real or effective UIDs or GIDs, can be assigned to rights profiles.

The following figure uses the Operator role and the Printer Management rights profiles as examples to demonstrate RBAC relationships.



**FIGURE 17-2** How Solaris RBAC Elements Relate

The Operator role is used to maintain printers and perform media backup. The user johnDoe is assigned to the Operator role and can assume it by supplying the Operator password.

The Operator rights profile has been assigned to the Operator role. The Operator rights profile has two supplementary profiles assigned to it, Printer Management and Media Backup, which reflect the Operator role's primary tasks.

The Printer Management rights profile is for managing printers, print daemons, and spoolers. Three authorizations are assigned to the Printer Management rights profile: `solaris.admin.printer.read`, `solaris.admin.printer.delete`, and `solaris.admin.printer.modify`. These authorizations allows users to manipulate information in the printer queue. The Printer Management profile also has a number of commands with security attributes that are assigned to it, such as `/usr/sbin/accept` with `euid=0` and `/usr/ucb/lpq` with `euid=lp`.

---

## Privileged Applications

Applications that can override system controls are considered *privileged applications*.



## Applications That Check UIDs and GIDs

Privileged applications that check for `root` or some other special UID or GID have long existed in UNIX. The RBAC rights profile mechanism enables you to specify the UID or GID for specific commands. Instead of changing the ID on a command that anyone can access, you can isolate the command with execution security attributes in the rights profile. A user or role with that rights profile can then run the program without having to become `root`.

IDs can be specified as real or effective. Assigning effective IDs is preferred over assigning real IDs. Effective IDs are equivalent to the `setuid` feature in the file permission bits and identify the UID for auditing. However, because some shell scripts and programs require a real UID of `root`, real IDs can be set as well. For example, the `pkgadd` command requires a real rather than an effective UID. If you encounter a command where the effective UID is not sufficient to run the command, you need to change the privilege to a real UID by using the Set Security Attributes option in the SMC Right Properties dialog boxes as described in “Creating or Changing a Rights Profile” on page 270.

## Applications That Check Authorizations

RBAC additionally provides commands that check authorizations. By definition, `root` has all authorizations and thus can run any application. Currently, the applications that check for authorizations include the following:

- The entire Solaris Management Console suite of tools
- The batch job-related commands – `at`, `atq`, `batch`, and `crontab`
- Device-oriented commands – `allocate`, `deallocate`, `list_devices`, and `cdrw`.

## Profile Shell

Authorized users can obtain privileged applications from the Solaris Management Console launcher or on the command line from a *profile shell*. A profile shell is a special kind of shell that enables access to the privileged applications that are assigned to the profile. Profile shells are launched when the user runs `su` to assume a role. The profile shells are `pfsh`, `pfcsh`, and `pfksh`. They correspond to Bourne shell (`sh`), C shell (`csh`), and Korn shell (`ksh`), respectively.

---

## RBAC Roles

A *role* is a special type of user account from which you can run privileged applications. Roles are created in the same general manner as user accounts, with a home directory, groups, password, and so on. The capabilities of a role are a function of the rights profiles and authorizations that are assigned to it. Roles do not have inheritance.

When a user assumes a role, the role's attributes replace all user attributes. Role information is stored in the `passwd`, `shadow`, `user_attr`, and `audit_user` databases. For detailed information on setting up roles, see "Configuring Recommended Roles" on page 279, "Creating Roles" on page 265, and "Changing Role Properties" on page 268.

All users who can assume the same role have the same role home directory, operate in the same environment, and have access to the same files. Users can assume roles from the command line by running `su` and supplying the role name and password. Users can also assume a role when they open a Solaris Management Console tool.

Users cannot log in directly to a role. For this reason, it is useful to make `root` a role to prevent anonymous `root` login. See "Making `Root` a Role" on page 261. Users must log in to their user account first. A user cannot assume a role directly from another role. A user's real UID can always be audited.

No predefined roles are shipped with the Solaris 9 software. As stated earlier in this chapter, you can easily configure the three recommended roles.

---

## RBAC Authorizations

An *authorization* is a discrete right that can be granted to a role or user. RBAC-compliant applications can check a user's authorizations prior to granting access to the application or specific operations within it. This check replaces the check in conventional UNIX applications for `UID=0`. For more information on authorizations, see "Authorizations" on page 284, "The `auth_attr` Database" on page 288, and "Commands That Require Authorizations" on page 293.

---

## RBAC Rights Profiles

A *rights profile* is a collection of system overrides that can be assigned to a role or user. A rights profile can contain commands with effective or real UIDs or GIDs defined, authorizations, and other rights profiles. Rights profile information is split between the `prof_attr` and `exec_attr` databases. For more information on rights profiles, see “Contents of Rights Profiles” on page 280, “The `prof_attr` Database” on page 290, and “The `exec_attr` Database” on page 291.

---

## Name Service Scope

Name service scope is an important concept for understanding RBAC. The scope in which a role can operate might apply to an individual host or to all hosts that are served by a name service such as NIS, NIS+, or LDAP. The precedence of local configuration files versus distributed databases is specified in the file `/etc/nsswitch.conf`. A lookup stops at the first match. For example, if a profile exists in two scopes, only the entries in the first scope are used.



---

## Role-Based Access Control (Tasks)

---

This chapter covers tasks that you can use to manage RBAC elements. The following is a list of the task maps in this chapter. To find the tasks for the initial setup of RBAC, see “Configuring RBAC (Task Map)” on page 254. For general management of the RBAC elements, see “Managing RBAC Information (Task Map)” on page 262.

The topics that are covered in this chapter include the following:

- “Planning for RBAC” on page 254
- “First-Time Use of the User Tool Collection” on page 256
- “Setting Up Initial Users” on page 257
- “Setting Up Initial Roles” on page 259
- “Making `Root` a Role” on page 261
- “Using Privileged Applications” on page 263
- “Creating Roles” on page 265
- “Changing Role Properties” on page 268
- “Creating or Changing a Rights Profile” on page 270
- “Modifying a User’s RBAC Properties” on page 274
- “Securing Legacy Applications” on page 276

The preferred method for performing RBAC-related tasks is through the Solaris Management Console. The console tools for managing the RBAC elements are contained in the User Tool Collection.

You can also operate on local files with the Solaris Management Console command-line interface and other command-line interfaces. The Solaris Management Console commands require authentication to connect to the server. As a result, they are not practical in scripts. The other commands require superuser or a role, and cannot be applied to databases in a name service.

---

**Tip** – Another drawback to using the command line to manage RBAC information is that the edit may not take effect immediately. To enable the edit, you need to stop and restart the name service cache daemon, `nscd(1M)`.

---

---

## Configuring RBAC (Task Map)

Task	Description	For Instructions
1. Plan for RBAC	Learn the concepts behind RBAC, examine your site's security needs, and plan how to integrate RBAC into your operation.	"How to Plan Your RBAC Implementation" on page 254
2. Start the User tools from the Solaris Management Console	All RBAC tasks can be performed by the User tools.	"How to Run the User Tool Collection" on page 256
3. Install initial users if needed	One or more existing users must be available for assignment to the first role.	"How to Create Initial Users by Using the User Accounts Tool" on page 257
4. Install the first role	The first role, typically Primary Administrator, needs to be installed by <code>root</code> user.	"How to Run the User Tool Collection" on page 256
5. (Optional) Make <code>root</code> a role	To eliminate anonymous <code>root</code> login, <code>root</code> can be made a role.	"How to Make <code>Root</code> a Role" on page 261

---

## Planning for RBAC

RBAC can be an integral part of how an organization manages its information resources. Planning requires a thorough knowledge of the RBAC capabilities as well as the security requirements of the organization.

### ▼ How to Plan Your RBAC Implementation

#### 1. Learn the basic RBAC concepts.

Read Chapter 17. Using RBAC to administer a system is very different from using conventional UNIX. You should be familiar with the RBAC concepts before you start your implementation. For greater detail, see Chapter 19.

#### 2. Examine your security policy.

Your organization's security policy should detail the potential threats to your system, measure the risk of each threat, and have a strategy to counter these threats. Isolating the security-relevant tasks through RBAC can be a part of the strategy. Although you

can install the suggested roles and their configurations as is, you might need to customize your RBAC configuration to adhere to your security policy.

### 3. Decide how much RBAC your organization needs.

Depending on your security needs, you can use varying degrees of RBAC, as follows:

- **No RBAC** – You can perform all tasks as `root` user. In this instance, you log in as yourself and when you select a console tool, you type `root` as the user.
- **Root as a Role** – This method eliminates anonymous `root` logins by preventing all users from logging in as `root`. Instead, they must log in as normal users prior to assuming the `root` role. See “Making `Root` a Role” on page 261.
- **Single Role Only** – This method adds the Primary Administrator role only and is similar to the superuser model.
- **Suggested Roles** – Three suggested roles that can be easily configured are available: Primary Administrator, System Administrator, and Operator. These roles are suitable for organizations with administrators at different levels of responsibility whose job capabilities fit the suggested roles.
- **Custom Roles** – You can create your own roles to meet the security requirements of your organization. The new roles can be based on existing or customized rights profiles.

### 4. Decide which suggested roles are appropriate for your organization.

Review the capabilities of the suggested roles and default rights profiles. Three rights profiles are available for configuring the suggested roles:

- **Primary Administrator rights profile** – For creating a role that can perform all administrative tasks, granting rights to others, and editing rights that are associated with administrative roles. A user in this role can assign the Primary Administrator role and the ability to grant rights to other users.
- **System Administrator rights profile** – For creating a role that can perform most nonsecurity administrative tasks. For example, the System Administrator can add new user accounts, but cannot set passwords or grant rights to other users.
- **Operator rights profile** – For creating a role that can perform simple administrative tasks, such as backup and restore, and printer maintenance.

These rights profiles enable administrators to configure the suggested roles by using a single rights profile instead of having to mix and match rights profiles.

To further examine rights profiles, use the Rights tool to display the contents. You can also refer to “Contents of Rights Profiles” on page 280 for a summary of some typical rights profiles. With the console tools, you can customize the roles and rights profiles that are provided to meet the needs of your organization.

### 5. Decide if any additional roles or rights profiles are appropriate for your organization.

Look for other applications or families of applications at your site that might benefit from restricted access. Applications that affect security, that can cause denial-of-service problems, or that require special administrator training are good candidates for RBAC.

- a. **Determine which commands are needed for the new task.**
  - b. **Decide which rights profile is appropriate for this task.**  
Check if an existing rights profile can handle this task or if a separate rights profile needs to be created.
  - c. **Determine which role is appropriate for this rights profile.**  
Decide if the rights profile for this task should be assigned to an existing role or if a new role should be created. If you use an existing role, check that the other rights profiles are appropriate for users who are assigned to this role.
6. **Decide which users should be assigned to the available roles.**  
According to the principle of least privilege, you should assign users to roles that are appropriate to their level of trust. Keeping users away from tasks that they do not need to use reduces potential problems.

---

## First-Time Use of the User Tool Collection

To install the initial users to their assigned roles, you first log in as yourself. When you authenticate yourself to the Solaris Management Console, specify `root` user.

### ▼ How to Run the User Tool Collection

1. **Log in as a normal user and start the Solaris Management Console.**

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

2. **Navigate to the User Tool Collection and click the icon, as follows:**

- a. **Find the icon that is labeled This Computer under Management Tools in the navigation pane.**
- b. **Click the turner icon to its left.**  
The turner icon is shaped like a lever. When the lever is horizontal, the contents of the folder are hidden. When the lever is vertical, the contents are displayed. Clicking the turner icon toggles the folder between the hidden and displayed states.
- c. **Click the turner icon next to the System Configuration folder to display its contents.**



d. Click the User icon to open the User Tool Collection.

The user login dialog box is displayed.

3. Type root and the root password in the Login: User Name dialog box. Click OK.

Generally, you should type your user name here and then assume a role. However, for the first time, you need to be root user because no roles exist yet. This step opens the User Tool Collection (see the following figure).




FIGURE 18–1 The User Tool Collection

---

## Setting Up Initial Users

If all users who are assigned to roles are already installed on this system, you can skip this task and go to “Setting Up Initial Roles” on page 259.

### ▼ How to Create Initial Users by Using the User Accounts Tool

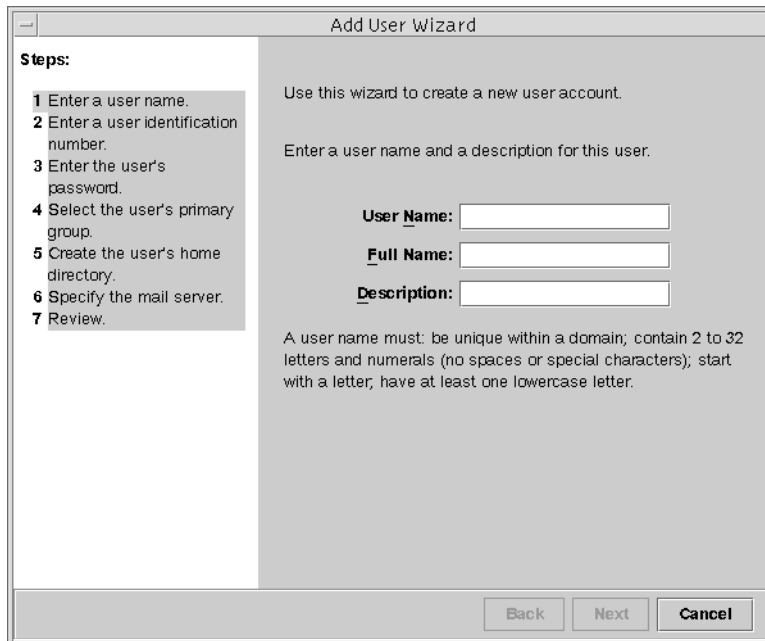
1. Click the User Accounts tool icon in either the navigation pane or the view pane of the User Tool Collection.

The User Accounts tool is started. The Action menu now provides options for this tool.

**2. Select Add User->With Wizard from the Action menu.**

This step starts the Add User wizard, a series of dialog boxes that request information that is necessary for configuring a user. Use the Next and Back buttons to navigate between dialog boxes. Note that the Next button does not become active until all required fields have been filled in. The last dialog box is for reviewing the entered data, at which point you can go back to change entries or click Finish to save the new role.

The following figure shows the first dialog box, Step 1: Enter a user name.



**FIGURE 18-2** Add User Wizard

**3. Type the name of the first user and the other identification information.**

**4. In the Step 2: Enter a User Identification Number dialog box, type the UID.**

This entry should match the existing UID for the user.

**5. In the Step 3: Enter the User's Password dialog box, indicate whether you or the user will be setting the password.**

If you are setting up this account for yourself, click the second option. Then, type and confirm your password.

**6. In the Step 4: Select the User's Primary Group dialog box, select the appropriate group.**

7. In the Step 5: Create the User's Home Directory dialog box, specify the path for the home directory.

8. In the Step 6: Specify the Mail Server dialog box, check out the default mail server and mailbox.

You can change these settings later in the User Properties dialog box.

9. Check the information in the Review dialog box. Click Finish to save, or click Back to reenter information.

If you discover missing or incorrect information, click the Back button repeatedly to display the dialog box where the incorrect information is displayed. Then, click Next repeatedly to return to the Review dialog box.

---

## Setting Up Initial Roles

The first role to create is the role that is responsible for managing users and roles, typically the Primary Administrator. First, you should install the users and the roles on your local host. After you have set up a toolbox for the name service scope, you need to create the same users and roles in the name service. See "Using the Solaris Management Tools in a Name Service Environment (Task Map)" in *System Administration Guide: Basic Administration*. After the first role is established and assigned to you, then you can run the console tools by assuming a role instead of becoming `root`.

### ▼ How to Create the First Role (Primary Administrator) by Using the Administrative Roles Tool

To install the first role, you should log in as yourself. When you authenticate yourself to the Solaris Management Console, specify `root` user. You should first install the role on your local host. After the first role is established and assigned to you, you can run the console tools by assuming a role instead of as `root` user.

1. Type `root` and the `root` password in the Login: User Name dialog box. Click OK.

2. Click the Administrative Roles icon in either the navigation pane or the view pane of the User Tool Collection.

The Administrative Roles tool is started. The Action menu now provides options for this tool.

3. Select Add Administrative Role from the Action menu.

This step starts the Add Administrative Role wizard, a series of dialog boxes that request information that is necessary for configuring a role. Use the Next and Back buttons to navigate between dialog boxes. Note that the Next button does not become active until all required fields have been filled in. The last dialog box is for reviewing the entered data, at which point you can go back to change entries or click Finish to save the new role.

The following figure shows the first dialog box, Step 1: Enter a Role Name.

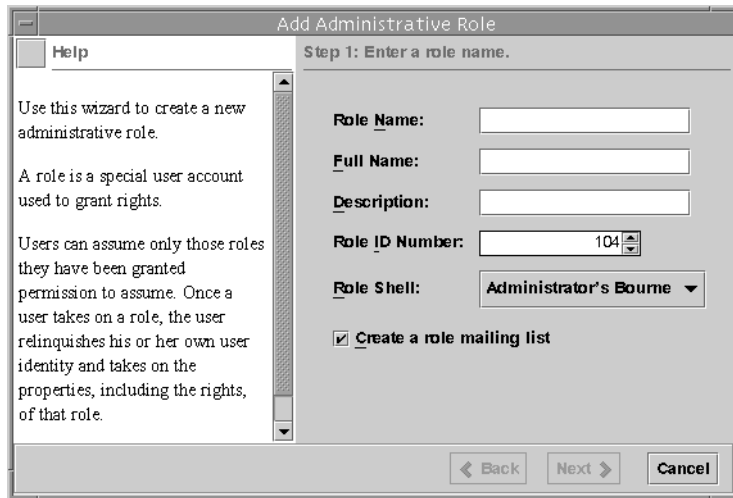


FIGURE 18–3 Add Administrative Role Wizard

**4. Type `primaryadmin`, or whatever role name you are using, and the other identification information.**

If you select the role mailing list option, you can create an alias of users who can assume this role.

**5. In the Step 2: Enter a Role Password dialog box, type the password for the new role in the Role Password field and again in the Confirm Password field.**

Confirmation helps prevent a misspelled password from being saved.

**6. In the Step 3: Enter Role Rights dialog box, select the Primary Administrator rights profile.**

Double-click the Primary Administrator rights profile in the Available Rights column (on the left). The rights profiles in the Granted Rights column (on the right) are the rights profiles that are assigned to this role. In this instance, only the Primary Administrator rights profile is needed.

**7. In the Step 4: Select a Home Directory dialog box, specify the server and path for the home directory.**

8. In the Step 5: Assign Users to This Role dialog box, type the login names for any users to be assigned to the Primary Administrator role.

Any users that you add must be defined in the same name service scope in which you are working. If you selected the role mailing list option in the Step 1: Enter a Role Name dialog box, these users will receive email that is addressed to the Primary Administrator role.

9. Check the information in the Review dialog box. Click Finish to save, or click Back to reenter information.

If you discover missing or incorrect information, click the Back button repeatedly to display the dialog box where the incorrect information is displayed. Then, click Next repeatedly to return to the Review dialog box.

10. Open a terminal window, become root, and start and stop the name service cache daemon.

The new role does not take effect until the name service cache daemon is restarted. After becoming root, type as follows:

```
# /etc/init.d/nscd stop
# /etc/init.d/nscd start
```

---

## Making Root a Role

This procedure shows how to change `root` from a user to a role within a local scope. Changing `root` to a role prevents users from logging in to that server directly as `root`. Users must first log in as themselves so their UIDs are available for auditing.



---

**Caution** – If you make `root` a role without assigning it to a valid user or without a currently existing role equivalent to `root`, no one can become `root`.

---

### ▼ How to Make Root a Role

1. Log in to the target server.
2. Become superuser.
3. Edit the `/etc/user_attr` file.

Here is an excerpt from a typical `user_attr` file.

```
root:::type=normal;auths=solaris.*,solaris.grant;profiles=All
johnDoe:::type=normal
```

4. Check that your name is in the file.

5. Add root to the roles that are assigned to your record.

Assign the root role to any applicable users. If you intend to use primaryadmin as your most powerful role, you do not have to assign root to any users.

```
johnDoe:::type=normal;roles=root
```

6. Go to the root record in the file and change type=normal to type=root.

```
root:::type=role;auths=solaris.*,solaris.grant;profiles=All
```

7. Save the file.

---

## Managing RBAC Information (Task Map)

The following task map shows where to obtain information for performing specific RBAC tasks.

Task	Description	For Instructions
Use privileged applications	To run applications that can affect security or system operations requires becoming superuser or assuming a role.	"How to Assume a Role in the Console Tools" on page 264 "How to Assume a Role at the Command Line" on page 263
Create roles	To add new roles, that is, special identities for running privileged applications.	"How to Create a Role by Using the Administrative Roles Tool" on page 265 "How to Create a Role From the Command Line" on page 267
Change role properties	To change the properties of a role, that is, the assigned users, rights profiles, and authorizations that are assigned to a role.	"How to Change a Role by Using the Administrative Roles Tool" on page 268 "How to Change a Role From the Command Line" on page 270

Task	Description	For Instructions
Create or change rights profiles	To add or change a rights profile, including the assignment of authorizations, commands with security attributes, and supplementary rights profiles.	<p>“How to Create or Change a Rights Profile by Using the Rights Tool” on page 270</p> <p>“How to Change Rights Profiles From the Command Line” on page 274</p>
Change a user’s RBAC properties	To change the roles, rights profiles, or authorizations that are assigned to a user.	<p>“How to Modify a User’s RBAC Properties by Using the User Accounts Tool” on page 275</p> <p>“How to Modify a User’s RBAC Properties From the Command Line” on page 275</p>
Secure legacy applications	To turn on the set ID permissions for legacy applications. Scripts can contain commands with set IDs. Legacy applications can check for authorizations, if appropriate.	<p>“How to Add Security Attributes to a Legacy Application” on page 276</p> <p>“How to Add Security Attributes to Commands in a Script” on page 276</p> <p>“How to Check for Authorizations in a Script or Program” on page 276</p>

These procedures manage the elements that are used in role-based access control (RBAC). For user management procedures, refer to “Managing User Accounts and Groups (Tasks)” in *System Administration Guide: Basic Administration*.

---

## Using Privileged Applications

To run privileged applications, you must first become superuser or assume a role. Although running privileged applications as a normal user is possible, it is discouraged to avoid errors that are caused by users who inadvertently exercise this privilege.

### ▼ How to Assume a Role at the Command Line

1. Use the `su` command as follows:

```
% su my-role
Password: my-role-password
#
```

Typing `su` by itself lets you become superuser. Typing `su` with a role name lets you assume that role (if it has been assigned to you). You must supply the appropriate password. Assuming a role switches the command line to the profile shell for that role. The profile shell has been modified to run commands with the security attributes that are assigned in the role's rights profiles.

**2. Type a command in the shell.**

The command is executed with any assigned security attributes and `setuid` or `setgid` permissions.

## ▼ How to Assume a Role in the Console Tools

**1. Start the Solaris Management Console.**

Use one of the following methods:

- Type `smc` at the command line.
- Click the Solaris Management Console icon in the Tools subpanel.
- Double-click the Solaris Management Console icon in the Application Manager.

All Solaris Management Console tools have extensive context-sensitive help that document each field. In addition, you can access various help topics from the Help menu. Note that it does not matter whether you are logged in as `root` or as a normal user when you start the console.

**2. Select the toolbox for your task.**

Navigate to the toolbox that contains the tool or collection in the appropriate scope and click the icon. The scopes are files (local), NIS, NIS+, and LDAP. If the appropriate toolbox is not displayed in the navigation pane, choose Open Toolbox from the Console menu and load the relevant toolbox.

**3. Select the tool.**

Navigate to the tool or collection to be used and click the icon. The tools for managing the RBAC elements are all part of the User Tool Collection.

**4. Authenticate yourself in the Login: User Name dialog box.**

Your choices are the following:

- Type your user name and password to assume a role or to operate as a normal user.
- Type `root` and the `root` password to operate as superuser.

Note that if you have not yet set up any roles or if the roles that are set up cannot perform the appropriate tasks, you need to log in as `root`. If you authenticate yourself as `root` (or as a user with no roles assigned), the tools are loaded into the console and you can proceed to Step 6.

**5. Authenticate yourself in the Login: Role dialog box.**



The Role option menu in the dialog box displays the roles that are assigned to you. Choose a role and type the role password. If you are to operate as a normal user, type your user name and password.

6. **Navigate to the tool to be run and click the icon.**

---

## Creating Roles

To create a role, you must either assume a role that has the Primary Administrator rights profile assigned to it or run as `root` user. See “RBAC Roles” on page 250 and “Configuring Recommended Roles” on page 279 to learn more about roles.

### ▼ How to Create a Role by Using the Administrative Roles Tool

1. **Start the Administrative Roles tool.**

Run the Administrative Roles tool, start the Solaris Management Console, as described in “How to Assume a Role in the Console Tools” on page 264. Then, open the User Tool Collection, and click the Administrative Roles icon.

2. **Start the Add Administrative Role wizard.**

Select Add Administrative Role from the Action menu to start the Add Administrative Role wizard for configuring roles.

3. **Fill in the fields in the series of dialog boxes. Click Finish when done.**

Use the Next and Back buttons to navigate between dialog boxes. Note that the Next button does not become active until all required fields have been filled in. The last dialog box is for reviewing the entered data, at which point you can go back to change entries or click Finish to save the new role. Table 18–1 summarizes the dialog boxes.

4. **Open a terminal window, become `root`, and start and stop the name service cache daemon.**

The new role does not take effect until the name service cache daemon is restarted. After becoming `root`, type as follows:

```
# /etc/init.d/nscd stop  
# /etc/init.d/nscd start
```

**TABLE 18-1** Add Administrative Role Wizard: Dialog Boxes and Fields

<b>Dialog Box</b>	<b>Fields</b>	<b>Field Description</b>
Step 1: Enter a role name	Role Name	Short name of the role.
	Full Name	Long version of the name.
	Description	Description of the role.
	Role ID Number	UID for the role, automatically incremented.
	Role Shell	The profile shells that are available to roles: Administrator's C, Administrator's Bourne, or Administrator's Korn shell.
	Create a role mailing list	Makes a mailing list for users who are assigned to this role.
Step 2: Enter a role password	Role Password	*****
	Confirm Password	*****
Step 3: Select role rights	Available Rights / Granted Rights	Assigns or removes a role's rights profiles.  Note that the system does not prevent you from typing multiple occurrences of the same command. The attributes that are assigned to the first occurrence of a command in a rights profile have precedence and all subsequent occurrences are ignored. Use the Up and Down arrows to change the order.
	Step 4: Select a home directory	Server
Step 5: Assign users to this role	Path	Home directory path.
	Add	Adds users who can assume this role. Must be in the same scope.
	Delete	Deletes users who are assigned to this role.

## ▼ How to Create a Role From the Command Line

### 1. Become superuser or assume a role that is capable of creating other roles.

### 2. Select a method for creating a role:

- For roles in the local scope, use the `roleadd` command to specify a new local role and its attributes.
- Alternatively, for roles in the local scope, edit the `user_attr` file to add a user with `type=role`.

This method is recommended for emergencies only, as it is easy to make mistakes while you are typing.

- For roles in a name service, use the `smrole` command to specify the new role and its attributes.

This command requires authentication by superuser or a role that is capable of creating other roles. You can apply the `smrole` to all name services. This command runs as a client of the Solaris Management Console server.

### 3. Start and stop the name service cache daemon.

New roles do not take effect until the name service cache daemon is restarted. As root, type as follows:

```
# /etc/init.d/nscd stop
# /etc/init.d/nscd start
```

#### EXAMPLE 18–1 Creating a Custom Operator Role by Using the `smrole` Command

The following sequence demonstrates how a role is created with the `smrole` command. In this example, a new version of the Operator role is created that has assigned to it the standard Operator rights profile and the Media Restore rights profile.

```
% su primaryadmin
# /usr/sadm/bin/smrole add -H myHost -- -c "Custom Operator" -n oper2 -a johnDoe \
-d /export/home/oper2 -F "Backup/Restore Operator" -p "Operator" -p "Media Restore"
Authenticating as user: primaryadmin

Type /? for help, pressing <enter> accepts the default denoted by [ ]
Please enter a string value for: password :: <type primaryadmin password>

Loading Tool: com.sun.admin.usermgr.cli.role.UserMgrRoleCli from myHost
Login to myHost as user primaryadmin was successful.
Download of com.sun.admin.usermgr.cli.role.UserMgrRoleCli from myHost was successful.

Type /? for help, pressing <enter> accepts the default denoted by [ ]
Please enter a string value for: password :: <type oper2 password>

# /etc/init.d/nscd stop
# /etc/init.d/nscd start
```

**EXAMPLE 18-1** Creating a Custom Operator Role by Using the `smrole` Command  
(Continued)

To view the newly created role (and any other roles), use `smrole` with the `list` subcommand, as follows:

```
# /usr/sadm/bin/smrole list --
Authenticating as user: primaryadmin

Type /? for help, pressing <enter> accepts the default denoted by [ ]
Please enter a string value for: password :: <type primaryadmin password>

Loading Tool: com.sun.admin.usermgr.cli.role.UserMgrRoleCli from myHost
Login to myHost as user primaryadmin was successful.
Download of com.sun.admin.usermgr.cli.role.UserMgrRoleCli from myHost was successful.
root          0          Super-User
primaryadmin  100         Most powerful role
sysadmin      101         Performs non-security admin tasks
oper2         102         Backup/Restore Operator
```

---

## Changing Role Properties

To change a role, you must either assume a role that has the Primary Administrator rights profile assigned to it, or run the User Tool Collection as `root` user if roles have not yet been set up.

### ▼ How to Change a Role by Using the Administrative Roles Tool

**1. Start the Administrative Roles tool.**

To run the Administrative Roles tool, you need to start the Solaris Management Console, as described in “How to Assume a Role in the Console Tools” on page 264. Then, open the User Tool Collection, and click the Administrative Roles icon.

After the Administrative Roles tool starts, the icons for the existing roles are displayed in the view pane.

**2. Click the role to be changed and select the appropriate item from the Action menu, as follows:**

- To change users who are assigned to a role, select **Assign Administrative Role**. The Assign Administrative Role dialog box is displayed. The Assign Administrative Role dialog box is a modified version of the Role Properties dialog box and has a **Users** tab only. Use the **Add** field to assign a user in the current scope to this role. Use the **Delete** field to remove a user’s role assignment. Click **OK**.

to save.

- To change rights that are assigned to a role, select Assign Rights to Role.

The Assign Rights to Role dialog box is displayed. The Assign Rights to Role dialog box is a modified version of the Role Properties dialog box and has a Rights tab only. Use the Available Rights and Granted Rights columns to add or remove rights profiles for the selected role. Click OK to save.

- To change any of the role's properties, select Properties (or simply double-click the role icon).

The Role Properties dialog box is displayed, which provides access to all role properties (see the following figure and table). Use the tabs to navigate to any information to be changed, make your changes, and click OK to save.

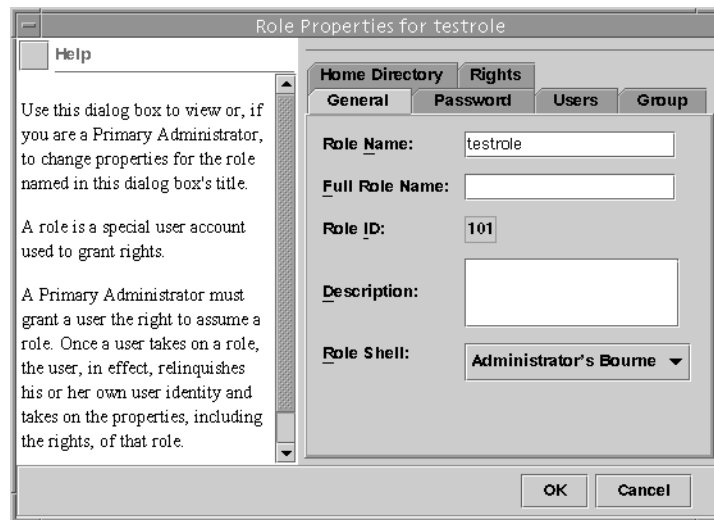


FIGURE 18-4 Role Properties Dialog Box

TABLE 18-2 Role Properties Summary

Tab	Tab Description
General	Specifies the role identification information and the default login shell.
Password	Specifies the role password.
Users	Specifies the users who are assigned to the role.
Group	Sets the role's primary groups and secondary groups for the purpose of accessing and creating files and directories.
Home Directory	Specifies the role's home directory, home directory server, automounting, and directory access.

**TABLE 18-2** Role Properties Summary (Continued)

Tab	Tab Description
Rights	Allows rights profiles to be assigned to the role. The precedence of the assigned rights profiles can be changed here.

## ▼ How to Change a Role From the Command Line

1. **Become superuser or assume a role that is capable of changing other roles.**
2. **Use the command that is appropriate for the task:**
  - Use the `rolemod` command to modify the attributes of a role that are defined locally.
  - Use the `roledel` command to delete a role that is defined locally.
  - Edit the `user_attr` file to change the authorizations or rights profiles that are assigned to a local role.

This method is recommended for emergencies only, as it is easy to make a mistake while you are typing.
  - Use the `smrole` command to modify the attributes of a role in a name service.

This command requires authentication as superuser or as a role that is capable of changing other roles. The `smrole` command runs as a client of the Solaris Management Console server.

---

## Creating or Changing a Rights Profile

To create or change a rights profile, you must either assume a role that has the Primary Administrator rights profile assigned to it, or run the User Tool Collection as `root` user if roles have not yet been set up. To learn more about rights profiles, see “RBAC Roles” on page 250 and “Configuring Recommended Roles” on page 279.

## ▼ How to Create or Change a Rights Profile by Using the Rights Tool

1. **Start the Rights tool.**

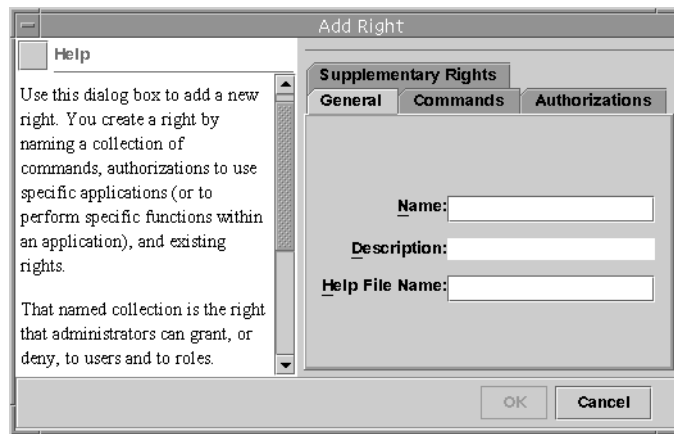
To run the Rights tool, you need to start the Solaris Management Console as described in “How to Assume a Role in the Console Tools” on page 264. Then, open the User Tool Collection, and click the Rights icon.

After the Rights tool starts, the icons for the existing rights profiles are displayed in the view pane.

**2. Take the appropriate action for creating or changing a rights profile:**

- To create a new rights profile, select Add Right from the Action menu.
- To change an existing rights profile, click the rights profile icon and select Properties from the Action menu (or simply double-click the rights profile icon).

Both actions display a version of the Rights Properties dialog box. The Add Right version (which follows) has a writable Name field. The standard Rights Properties dialog box has a read-only Name field because the name of a rights profile cannot be changed after it has been defined.



**FIGURE 18–5** Add Right Dialog Box

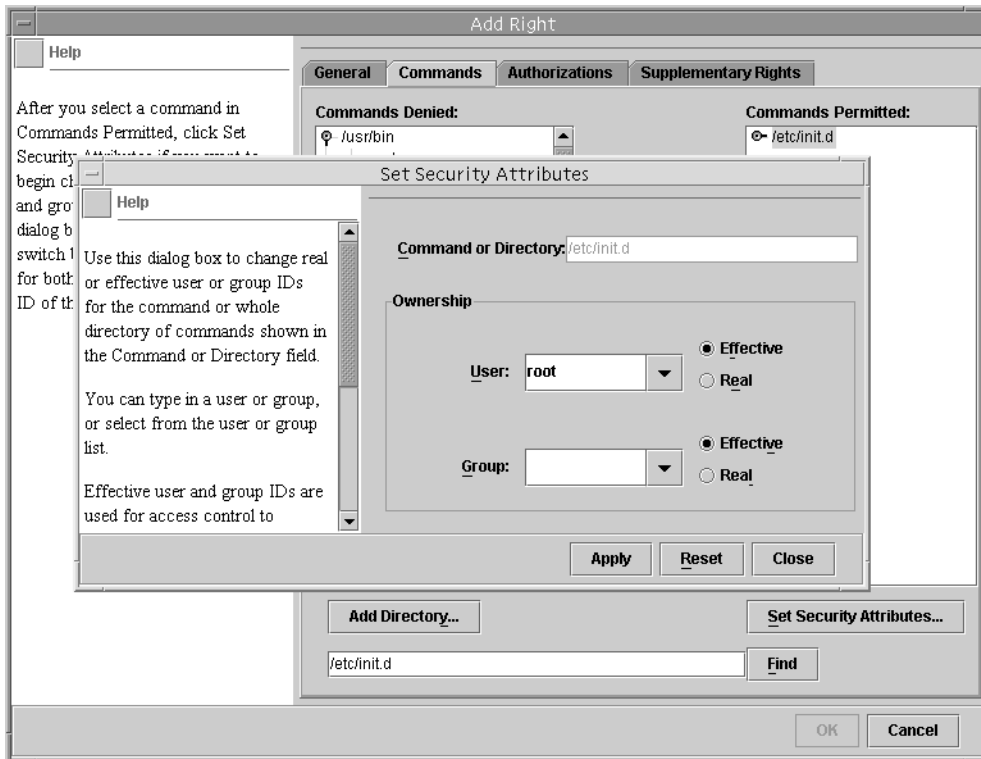
**3. Type the new information. Click OK to save the rights profile.**

The following table lists the tabs and fields in the Right Properties dialog box.

Tab	Field	Field Description
General	Name	Name of the new rights profile.
	Description	Description of the new rights profile.
	Help File Name	Name of the HTML help file for the new rights profile.
Commands	Add Directory	Opens a dialog box for adding directories that are not already in the Commands Denied or Commands Permitted columns.

Tab	Field	Field Description
	Commands Denied / Commands Permitted	Assigns or removes a rights profile's commands.
	Set Security Attributes	Opens a dialog box for assigning or removing a command's security attributes, that is, real or effective UIDs or GIDs (see Figure 18-6).  <b>Note</b> – Assigning effective IDs is preferred over assigning real IDs. Use real IDs only when they are required by the command, such as <code>pkgadd</code> .
	Find (command)	Searches the two command lists for the specified string.
Authorizations	Authorizations Excluded / Authorizations Included	Assigns or removes a rights profile's authorizations.
Supplementary Rights	Rights Excluded / Rights Included	Assigns or removes a rights profile's supplementary rights profiles.





**FIGURE 18-6** Adding Security Attributes to Commands

**EXAMPLE 18-2** Creating a New Rights Profile With the Rights Tool

The data in the following table shows sample data for a hypothetical rights profile that is called “Restart” could be created. The example rights profile, Restart, has the commands in the subdirectory `/etc/init.d` assigned to it. These commands have an effective UID of 0. This rights profile would be useful for administrators who are permitted to stop and start the daemons in `/etc/init.d`.

Tab	Field	Example
General	Name	Restart
	Description	For starting and stopping daemons in <code>/etc/init.d</code>
	Help File Name	<code>Restart.html</code>

**EXAMPLE 18-2** Creating a New Rights Profile With the Rights Tool (Continued)

Tab	Field	Example
Commands	Add Directory	Click Add Directory, type <code>/etc/init.d</code> in the dialog box, and click OK.
	Commands Denied / Commands Permitted	Select <code>/etc/init.d</code> and click Add to move the command to the Commands Permitted column.
	Set Security Attributes	Select <code>/etc/init.d</code> , click Set Security Attributes, and set Effective UID = root (see Figure 18-6).
	Find (command)	
Authorizations	Authorizations Excluded / Authorizations Included	
Supplementary Rights	Rights Excluded / Rights Included	

## ▼ How to Change Rights Profiles From the Command Line

1. **Become superuser or assume a role with the PrimaryAdministration rights profile.**
2. **Use the subcommand of `smprofile` that is appropriate for the task.**

This command requires authentication. You can apply the command to all name services. `smprofile` runs as a client of the Solaris Management Console server.

- To add a new profile, use `smprofile` with the `add` subcommand.
- To change an existing profile, use `smprofile` with the `modify` subcommand.

---

## Modifying a User's RBAC Properties

To modify a user's properties, you must either be running the User Tool Collection as root user or assume a role that has the Primary Administrator rights profile assigned to it.

## ▼ How to Modify a User's RBAC Properties by Using the User Accounts Tool

### 1. Start the User Accounts tool.

To run the User Accounts tool, you need to start the Solaris Management Console, as described in “How to Assume a Role in the Console Tools” on page 264. Then, open the User Tool Collection, and click the User Accounts icon.

After the User Accounts tool starts, the icons for the existing user accounts are displayed in the view pane.

### 2. Click the user account icon to be changed and select Properties from the Action menu (or simply double-click the user account icon).

### 3. Click the appropriate tab in the dialog box for the property to be changed, as follows:

- To change the roles that are assigned to the user, click the Roles tab and move the role assignment to be changed to the appropriate column: Available Roles or Assigned Roles.
- To change the rights profiles that are assigned to the user, click the Rights tab and move it to the appropriate column: Available Rights or Assigned Rights.

---

**Note** – It is not good practice to assign rights profiles directly to users. The preferred approach is to force users to assume roles in order to perform privileged applications. This strategy avoids the possibility of normal users abusing privileges.

---

## ▼ How to Modify a User's RBAC Properties From the Command Line

### 1. Become superuser or assume a role that can modify user files.

### 2. Use the appropriate command:

- To change the authorizations, roles, or rights profiles that are assigned to a user who is defined in the local scope, use the `usermod` command.
- Alternatively, to change the authorizations, roles, or rights profiles that are assigned to a user who is defined in the local scope, edit the `user_attr` file.  
This method is recommended for emergencies only, as it is easy to make a mistake while you are typing.
- To change the authorizations, roles, or rights profiles that are assigned to a user who is defined in a name service, use the `smuser` command.

This command requires authentication as superuser or as a role that is capable of changing user files. You can apply `smuser` to all name services. `smuser` runs as a client of the Solaris Management Console server.

---

## Securing Legacy Applications

This section discusses how to make legacy applications more secure. To add legacy applications to the Solaris Management Console, see “Adding Tools to the Solaris Management Console” in *System Administration Guide: Basic Administration*.

### ▼ How to Add Security Attributes to a Legacy Application

You add security attributes to a legacy application in the same way as you would for any command. You need to add the command (or its directory) to the Commands Denied column in the Commands tab of the Rights property dialog box. Then move the command to the Commands Permitted column.

### ▼ How to Add Security Attributes to Commands in a Script

If a command in a script needs to have the setUID bit set to run, simply add the security attributes to that command in the same rights profile. See “How to Create or Change a Rights Profile by Using the Rights Tool” on page 270.

### ▼ How to Check for Authorizations in a Script or Program

To have a script for authorizations, you need to add a test that is based on the `auths` command (see the `auths(1)` man page). For example, the following line would test if the user has the authorization entered as the `$1` argument:

```
if [ `'/usr/bin/auths|/usr/xpg4/bin/grep $1` ]; then
    echo Auth granted
else
    echo Auth denied
fi
```

To be more complete, the test should include logic that checks for other authorizations that use wildcards. For example, to test if the user has the `solaris.admin.usermgr.write` authorization, you need to check for the strings: `solaris.admin.usermgr.write`, `solaris.admin.usermgr.*`, `solaris.admin.*`, and `solaris.*`.

If you are writing a program, use the `getauthattr()` function to test for the authorization.



## Role-Based Access Control (Reference)

---

This chapter provides additional information that supplements Chapter 17.

This is a list of the reference information in this chapter:

- “Configuring Recommended Roles” on page 279
- “Contents of Rights Profiles” on page 280
- “Authorizations” on page 284
- “RBAC Database Relationships” on page 286
- “The user\_attr Database” on page 287
- “The auth\_attr Database” on page 288
- “The prof\_attr Database” on page 290
- “The exec\_attr Database” on page 291
- “Command-Line Applications for Managing RBAC” on page 292
- “Commands That Require Authorizations” on page 293

For information on RBAC tasks, see Chapter 18.

---

## RBAC Elements: Reference Information

This section describes the role-based access control (RBAC) elements in detail.

### Configuring Recommended Roles

No predefined roles are shipped with the Solaris 9 software. Management at a customer site must decide what types of roles should be set up. However, three recommended roles can be readily configured by assigning the appropriate predefined rights profile to the corresponding roles:

- **Primary Administrator rights profile** – For creating a role that can perform all administrative tasks, granting rights to others, and editing rights that are associated with administrative roles. A user in this role can assign the Primary Administrator role and the ability to grant rights to other users.
- **System Administrator rights profile** – For creating a role that can perform most nonsecurity administrative tasks. For example, the System Administrator can add new user accounts, but cannot set passwords or grant rights to other users.
- **Operator rights profile** – For creating a role that can perform simple administrative tasks, such as backup and restore, and printer maintenance.

These rights profiles enable administrators to configure the suggested roles by using a single rights profile instead of having to mix and match rights profiles.

Those sites that customize roles should closely check the order of the rights profiles that are assigned to the role. The system does not prevent someone from typing multiple occurrences of the same command. The attributes that are assigned to the first occurrence of a command in a rights profile take precedence and all subsequent occurrences are ignored.

---

**Note** – You can also set up `root` as a role through a manual process. This method prevents users from logging in directly as `root`, forcing them to log in as themselves first. See “Making Root a Role” on page 261.

---

## Contents of Rights Profiles

This section describes some typical rights profiles.

- The All rights profile provides role access to commands without security attributes.
- The Primary Administrator rights profile is designed specifically for the Primary Administrator role. The Primary Administrator rights profile allows the use of wildcards.
- The System Administrator rights profile is designed specifically for the System Administrator role. The System Administrator rights profile uses discrete supplementary profiles to create a powerful role.
- The Operator rights profile is designed specifically for the Operator role. The Operator rights profile uses a few discrete supplementary profiles to create a simple role.
- The Basic Solaris User rights profile shows how the `policy.conf` file can be used to assign tasks to users that are not related to security.
- The Printer Management rights profile exemplifies a profile that is dedicated to a single area of administration.



The tables in the following sections show the purpose and the contents of these rights profiles, including the commands, authorizations, supplementary rights, rights profiles, and associated help files.

Help files are in HTML and can be readily customized, if required. These files reside in the `/usr/lib/help/auths/locale/C` directory.

The Solaris Management Console Rights tool provides another way of inspecting the contents of the rights profiles.

## All Rights Profile

The All rights profile uses the wildcard to include all commands, except for those commands without security attributes. This rights profile provides a role with access to all commands that are not explicitly assigned in other rights profiles. Without the All rights profile or some other rights profiles that use wildcards, a role has access to explicitly assigned commands only, which is not very practical.

Because commands in rights profiles are interpreted in the order in which they occur, any wildcard settings should be positioned last so that explicit attribute assignments are not inadvertently overridden. The All rights profile, if used, should be the final rights profile that is assigned.

**TABLE 19-1** Contents of All Rights Profile

Purpose	Contents
To execute any command as the user or role	<b>Commands:</b> * <b>Help File:</b> RtAll.html

## Primary Administrator Rights Profile

The Primary Administrator rights profile is assigned the most powerful role on the system, effectively providing that role with superuser capabilities.

- The `solaris.*` authorization effectively assigns all of the authorizations that are provided by the Solaris software.
- The `solaris.grant` authorization lets a role assign any authorization to any rights profile, role, or user.
- The command assignment `*:uid=0;gid=0` provides the ability to run any command with UID=0 and GID=0.

The help file `RtPriAdmin.html` is identified so that a site can modify it if necessary. Help files are stored in the `/usr/lib/help/auths/locale/C` directory.

Note also that if the Primary Administrator rights profile is not consistent with a site's security policy, it can be modified or not assigned at all. However, the security capabilities in the Primary Administrator rights profile would need to be handled in one or more other rights profiles.

**TABLE 19-2** Contents of Primary Administrator Rights Profile

Purpose	Contents
To perform all administrative tasks	<b>Commands:</b> *  <b>Authorizations:</b> solaris.*, solaris.grant  <b>Help File:</b> RtPriAdmin.html

## System Administrator Rights Profile

The System Administrator rights profile is intended for the System Administrator role. Because the System Administrator does not have the broad powers of the Primary Administrator, no wildcards are used. Instead, discrete administrative rights profiles that do not deal with security are assigned. The commands that are assigned to the supplementary rights profiles are not shown in the following table.

Notice that the All rights profile is assigned at the end of the list of supplementary rights profiles.

**TABLE 19-3** Contents of System Administrator Rights Profile

Purpose	Contents
To perform most nonsecurity administrative tasks	<b>Supplementary rights profiles:</b> Audit Review, Printer Management, Cron Management, Device Management, File System Management, Mail Management, Maintenance and Repair, Media Backup, Media Restore, Name Service Management, Network Management, Object Access Management, Process Management, Software Installation, User Management, All  <b>Help File:</b> RtSysAdmin.html

## Operator Rights Profile

The Operator rights profile is a less powerful administrative rights profile that provides the ability to do backups and printer maintenance. The ability to restore files has more security consequences, and the default is not to assign it to this rights profile.

**TABLE 19-4** Contents of Operator Rights Profile

<b>Purpose</b>	<b>Contents</b>
To perform simple administrative tasks	<b>Supplementary rights profiles:</b> Printer Management, Media Backup, All  <b>Help File:</b> RtOperator.html

## Basic Solaris User Rights Profile for User

By default, the Basic Solaris User rights profile is assigned automatically to all users through the `policy.conf` file. This rights profile provides basic authorizations that are useful in normal operations. Note that the convenience that is offered by the Basic Solaris User rights profile must be balanced against site security requirements. Those sites that need stricter security might prefer to remove this rights profile from the `policy.conf` file.

**TABLE 19-5** Contents of Basic Solaris User Rights Profile

<b>Purpose</b>	<b>Contents</b>
To automatically assign rights to all users	<b>Authorizations:</b> <code>solaris.profmgr.read</code> , <code>solaris.admin.usermgr.read</code> , <code>solaris.admin.logsvc.read</code> , <code>solaris.admin.fsmgr.read</code> , <code>solaris.admin.serialmgr.read</code> , <code>solaris.admin.diskmgr.read</code> , <code>solaris.admin.procmgr.user</code> , <code>solaris.compsys.read</code> , <code>solaris.admin.printer.read</code> , <code>solaris.admin.prodreg.read</code> , <code>solaris.admin.dcmgr.read</code>  <b>Supplementary rights profiles:</b> All  <b>Help File:</b> RtDefault.html

## Printer Management Rights Profile

Printer Management is a typical rights profile that is intended for a specific task area. Both authorizations and commands are assigned to the Printer Management rights profile. The following table shows a partial list of commands.

TABLE 19-6 Contents of Printer Management Rights Profile

Purpose	Contents
To manage printers, daemons, and spooling	<p><b>Authorizations:</b> <code>solaris.admin.printer.delete</code>, <code>solaris.admin.printer.modify</code>, <code>solaris.admin.printer.read</code></p> <p><b>Commands:</b> <code>/usr/sbin/accept:uid=lp</code>, <code>/usr/ucb/lpq:uid=0</code>, <code>/etc/init.d/lp:uid=0</code>, <code>/usr/bin/lpstat:uid=0</code>, <code>/usr/lib/lp/lpsched:uid=0</code>, <code>/usr/sbin/lpfilter:uid=lp</code></p> <p><b>Help File:</b> <code>RtPrntMngmnt.html</code></p>

## Authorizations

An *authorization* is a discrete right that can be granted to a role or user. Authorizations are checked by RBAC-compliant applications before a user gets access to the application or specific operations within it. This check replaces the tests in conventional UNIX applications for `UID=0`.

## Authorization Naming Convention

An authorization has a name that is used internally and in files (for example, `solaris.admin.usermgr.pswd`), and a short description, which appears in the graphical user interfaces (for example, `Change Passwords`).

By convention, authorization names consist of the reverse order of the Internet name of the supplier, the subject area, any subareas, and the function, which are all separated by dots. An example would be `com.xyzcorp.device.access`. Exceptions to this convention are the authorizations from Sun Microsystems, Inc., which use the prefix `solaris` instead of an Internet name. Sun's convention enables administrators to apply authorizations in a hierarchical fashion by using a wildcard (\*) to represent any strings to the right of a dot.

## Example of Authorization Granularity

As an example of how authorizations are used, consider the following. A user in the Operator role might be limited to the `solaris.admin.usermgr.read` authorization, which provides read but not write access to users' configuration files. The System Administrator role naturally has the `solaris.admin.usermgr.read` and also the `solaris.admin.usermgr.write` authorization for making changes to users' files. However, without the `solaris.admin.usermgr.pswd` authorization, the System Administrator cannot change passwords. The Primary Administrator has all three of these authorizations.

The `solaris.admin.usermgr.pswd` authorization is required to make password changes in the Solaris Management Console User Tool. This authorization is also required for using the password modification options in the `smuser`, `smmultiuser`, and `smrole` commands.

## Delegating Authorizations

An authorization that ends with the suffix `grant` permits a user or role to delegate to other users any assigned authorizations that begin with the same prefix.

For example, a role with the authorizations `solaris.admin.usermgr.grant` and `solaris.admin.usermgr.read` can delegate the `solaris.admin.usermgr.read` authorization to another user. A role with the `solaris.admin.usermgr.grant` and `solaris.admin.usermgr.*` can delegate any of the authorizations with the `solaris.admin.usermgr` prefix to other users.

---

## Databases That Support RBAC

The following four databases store the data for the RBAC elements:

- `user_attr` (extended user attributes database) – Associates users and roles with authorizations and rights
- `auth_attr` (authorization attributes database) – Defines authorizations and their attributes, and identifies the associated help file
- `prof_attr` (rights profile attributes database) – Defines rights profiles, lists the rights profile's assigned authorizations, and identifies the associated help file
- `exec_attr` (execution attributes database) – Identifies the commands with security attributes that are assigned to specific rights profiles

---

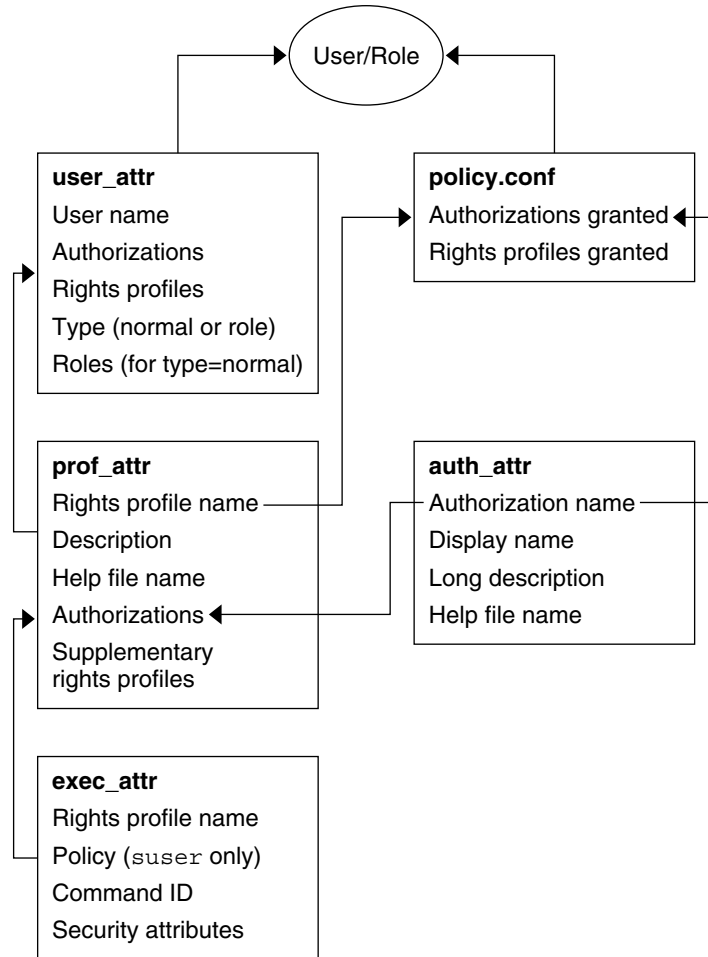
**Note** – The commands can also indicate a security policy. Currently, the only security policy that is available for the Solaris operating environment is `suser` (short for superuser). The `suser` policy is the default and it accommodates both the ID attributes and authorizations. The Trusted Solaris environment, which can interoperate with the Solaris environment, uses a policy called `tsol`. Additional policies might be available in future releases.

---

The `policy.conf` database is also important to the RBAC implementation. This database can contain authorizations and rights profiles that are to be applied by default to all users.

# RBAC Database Relationships

The following figure illustrates how the RBAC databases work together.



**FIGURE 19-1** RBAC Database Relations

The `user_attr` database stores the basic definitions for both users and roles, which are differentiated by the `type` field. The `user_attr` database contains the attributes that are shown in the figure, which includes a comma-separated list of rights profile names. The definitions of the rights profiles are split between two databases. The `prof_attr` database contains rights profile identification information, authorizations that are assigned to the profile, and supplementary profiles. The `exec_attr` database identifies the security policy and contains commands with their associated security

attributes. The `auth_attr` database supplies authorization information to the Sun Management Console tools. The `policy.conf` database supplies default authorizations and rights profiles that are to be applied to all users.

Each database uses a `key=value` syntax for storing attributes. This method accommodates future expansion of the databases and enables a system to continue if it encounters a key that is unknown to its policy.

The scope of the RBAC databases can apply to individual hosts or to all hosts that are served by a name service such as NIS, NIS+, or LDAP. The precedence of local configuration files versus distributed databases for the `user_attr` database is set by the precedence that is specified for the `passwd` entry in the file `/etc/nsswitch.conf`. The precedence for the `prof_attr` and `auth_attr` databases are individually set in `/etc/nsswitch.conf`. The `exec_attr` database uses the same precedence as `prof_attr`. For example, if a command with security attributes is assigned to a profile that exists in two scopes, only the entry in the first scope is used.

The databases can reside on a local system or can be administered by the NIS, NIS+, or LDAP name service.

You can edit the databases manually or manipulate them with the commands that are described in “Command-Line Applications for Managing RBAC” on page 292.

## The `user_attr` Database

The `user_attr` database contains user and role information that supplements the `passwd` and `shadow` databases. The `user_attr` database contains extended user attributes such as authorizations, rights profiles, and assigned roles. The fields in the `user_attr` database are separated by colons, as follows:

```
user:qualifier:res1:res2:attr
```

The following table describes these fields.

Field Name	Description
<code>user</code>	The name of the user or role as specified in the <code>passwd</code> database.
<code>qualifier</code>	Reserved for future use.
<code>res1</code>	Reserved for future use.
<code>res2</code>	Reserved for future use.

Field Name	Description
attr	<p>An optional list of semicolon-separated (;) key-value pairs that describes the security attributes to be applied when the user runs commands. The four valid keys are <code>type</code>, <code>auths</code>, <code>profiles</code>, and <code>roles</code>.</p> <ul style="list-style-type: none"> <li>■ The <code>type</code> key can be set to <code>normal</code>, if this account is for a normal user, or to <code>role</code>, if this account is for a role.</li> <li>■ The <code>auths</code> key specifies a comma-separated list of authorization names that are chosen from names that are defined in the <code>auth_attr</code> database. Authorization names can include the asterisk (*) character as a wildcard. For example, <code>solaris.device.*</code> means all of the Solaris device authorizations.</li> <li>■ The <code>profiles</code> key specifies an ordered, comma-separated list of rights profile names from the <code>prof_attr</code> database. The order of rights profiles works similarly to UNIX search paths. The first rights profile in the list that contains the command to be executed defines which (if any) attributes are to be applied to the command.</li> <li>■ The <code>roles</code> key can be assigned to the user through a comma-separated list of role names. Note that roles are defined in the same <code>user_attr</code> database. Roles are indicated by setting the <code>type</code> value to <code>role</code>. Roles cannot be assigned to other roles.</li> </ul>

The following example demonstrates how the Operator role is defined in a typical `user_attr` database and how it is assigned to user `johnDoe`. Roles and users are differentiated by the `type` keyword.

```
% grep operator /etc/user_attr
johnDoe:::type=normal;roles=sysadmin,operator
operator:::profiles=Operator;type=role
```

## The auth\_attr Database

All authorizations are stored in the `auth_attr` database. Authorizations can be assigned directly to users (or roles) in the `user_attr` database. Authorizations can also be assigned to rights profiles, which are assigned to users.

The fields in the `auth_attr` database are separated by colons, as follows:

```
authname:res1:res2:short_desc:long_desc:attr
```

The following table describes these fields.



Field Name	Description
authname	<p>A unique character string that is used to identify the authorization in the format <i>prefix.[suffix]</i>. Authorizations for the Solaris operating environment use <code>solaris</code> as a prefix. All other authorizations should use a prefix that begins with the reverse-order Internet domain name of the organization that creates the authorization (for example, <code>com.xyzcompany</code>). The suffix indicates what is being authorized, which is typically the functional area and operation.</p> <p>When the authname consists of a prefix and functional area and ends with a period, the authname serves as a heading to be used by applications in their GUIs, rather than as an actual authorization. The authname of <code>solaris.printmgr.</code> is an example of a heading.</p> <p>When authname ends with the word “grant,” the authname serves as a grant authorization and lets the user delegate authorizations with the same prefix and functional area to other users. The authname of <code>solaris.printmgr.grant</code> is an example of a grant authorization. <code>solaris.printmgr.grant</code> gives the user the right to delegate such authorizations as <code>solaris.printmgr.admin</code> and <code>solaris.printmgr.nobanner</code> to other users.</p>
res1	Reserved for future use.
res2	Reserved for future use.
short_desc	A terse name for the authorization that is suitable for display in user interfaces, such as in a scrolling list in a GUI.
long_desc	A long description. This field identifies the purpose of the authorization, the applications in which it is used, and the type of user who might be interested in using it. The long description can be displayed in the help text of an application.
attr	<p>An optional list of semicolon-separated (;) key-value pairs that describe the attributes of an authorization. Zero or more keys can be specified.</p> <p>The keyword <code>help</code> identifies a help file in HTML. Help files can be accessed from the <code>index.html</code> file in the <code>/usr/lib/help/auths/locale/C</code> directory.</p>

The following example shows an `auth_attr` database with some typical values.

```
% grep printer /etc/security/auth_attr
solaris.admin.printer.:Printer Information::help=AuthPrinterHeader.html
solaris.admin.printer.delete::Delete Printer Information::help=AuthPrinterDelete.html
solaris.admin.printer.modify::Update Printer Information::help=AuthPrinterModify.html
solaris.admin.printer.read::View Printer Information::help=AuthPrinterRead.html
```

Note that `solaris.admin.printer.` is defined to be a heading, because it ends in a dot (.). Headings are used by the GUIs to organize families of authorizations.

## The prof\_attr Database

The prof\_attr database stores the name, description, help file location, and authorizations that are assigned to rights profiles. The commands and security attributes that are assigned to rights profiles are stored in the exec\_attr database (see “The exec\_attr Database” on page 291). The fields in the prof\_attr database are separated by colons:

```
profname:res1:res2:desc:attr
```

The following table describes these fields.

Field Name	Description
profname	The name of the rights profile. Rights profile names are case-sensitive. This name is also used by the user_attr database to indicate rights profiles that are assigned to roles and users.
res1	Reserved for future use.
res2	Reserved for future use.
desc	A long description. This field should explain the purpose of the rights profile, including what type of user would be interested in using it. The long description should be suitable for display in the help text of an application.
attr	<p>An optional list of key-value pairs that are separated by semicolons (;) that describes the security attributes to apply to the object on execution. Zero or more keys can be specified. The two valid keys are help and auths.</p> <p>The keyword help identifies a help file in HTML. Help files can be accessed from the index.html file in the /usr/lib/help/auths/locale/C directory.</p> <p>The keyword auths specifies a comma-separated list of authorization names that are chosen from those names that are defined in the auth_attr database. Authorization names can be specified with the asterisk (*) character as a wildcard.</p>

The following example shows a typical prof\_attr database. Note that the Printer Management rights profile is a supplementary rights profile that is assigned to the Operator rights profile.

```
% grep 'Printer Management' /etc/security/prof_attr
Printer Management::Manage printers, daemons, spooling:help=RtPrntAdmin.html; \
auths=solaris.admin.printer.read,solaris.admin.printer.modify,solaris.admin.printer.delete \
Operator::Can perform simple administrative tasks:profiles=Printer Management,\
Media Backup,All;help=RtOperator.html
...
```

# The `exec_attr` Database

An execution attribute is a command that is associated with a specific UID or GID and that is assigned to a rights profile. The command with its security attributes can be run by users or roles to whom the rights profile is assigned.

The `exec_attr` database stores the definitions of the execution attributes.

The fields in the `exec_attr` database are separated by colons:

```
name:policy:type:res1:res2:id:attr
```

The following table describes these fields.

Field Name	Description
<code>name</code>	The name of the rights profile. Rights profile names are case-sensitive. The name refers to a rights profile in the <code>prof_attr</code> database.
<code>policy</code>	The security policy that is associated with this entry. Currently, <code>suser</code> (the superuser policy model) is the only valid entry.
<code>type</code>	The type of entity that is specified. Currently, the only valid entity type is <code>cmd</code> (command).
<code>res1</code>	Reserved for future use.
<code>res2</code>	Reserved for future use.
<code>id</code>	A string that identifies the entity. Commands should have the full path or a path with a wildcard. To specify arguments, write a script with the arguments and point the <code>id</code> to the script.
<code>attr</code>	<p>An optional list of semicolon (;) separated key-value pairs that describes the security attributes to apply to the entity on execution. Zero or more keys can be specified. The list of valid keywords depends on the policy that is enforced. The four valid keys are <code>eid</code>, <code>uid</code>, <code>egid</code>, and <code>gid</code>.</p> <p>The <code>eid</code> and <code>uid</code> keywords contain a single user name or a numeric user ID (UID). Commands that are designated with <code>eid</code> run with the effective UID indicated, which is similar to setting the <code>setuid</code> bit on an executable file. Commands that are designated with <code>uid</code> run with both the real and effective UIDs.</p> <p>The <code>egid</code> and <code>gid</code> keywords contain a single group name or numeric group ID (GID). Commands that are designated with <code>egid</code> run with the effective GID indicated, which is similar to setting the <code>setgid</code> bit on an executable file. Commands that are designated with <code>gid</code> run with both the real and effective GIDs.</p>

The following example shows some typical values from an `exec_attr` database.

```
% grep 'Printer Management' /etc/security/exec_attr
Printer Management:suser:cmd::/usr/sbin/accept:eid=lp
Printer Management:suser:cmd::/usr/ucb/lpq:eid=0
Printer Management:suser:cmd::/etc/init.d/lp:eid=0
```

## The `policy.conf` File

The `policy.conf` file provides a way of granting specific rights profiles and authorizations to all users. The two types of entries in the file consist of key-value pairs. They are the following:

- `AUTHS_GRANTED=authorizations` – Refers to one or more authorizations
- `PROFS_GRANTED=right profiles` – Refers to one or more rights profiles

The following example shows some typical values from a `policy.conf` database.

```
# grep AUTHS /etc/security/policy
AUTHS_GRANTED=solaris.device.cdrw

# grep PROFS /etc/security/policy
PROFS_GRANTED=Basic Solaris User
```

---

## RBAC Commands

This section lists commands that are used to administer RBAC. Also provided is a table of commands whose access can be controlled by authorizations.

## Command-Line Applications for Managing RBAC

In addition to editing the RBAC databases directly, the following commands are available for managing access to tasks with RBAC.

**TABLE 19-7** RBAC Administration Commands

Command	Description
<code>auths(1)</code>	Displays authorizations for a user.
<code>makedbm(1M)</code>	Makes a dbm file.
<code>nscd(1M)</code>	Name service cache daemon, useful for caching the <code>user_attr</code> , <code>prof_attr</code> , and <code>exec_attr</code> databases.

**TABLE 19-7** RBAC Administration Commands (Continued)

Command	Description
<code>pam_roles(5)</code>	Role account management module for PAM. Checks for the authorization to assume role.
<code>pfexec(1)</code>	Used by profile shells to execute commands with attributes that are specified in the <code>exec_attr</code> database.
<code>policy.conf(4)</code>	Configuration file for security policy. Lists granted authorizations.
<code>profiles(1)</code>	Displays rights profiles for a specified user.
<code>roles(1)</code>	Displays roles that are granted to a user.
<code>roleadd(1M)</code>	Adds a role to the system.
<code>roledel(1M)</code>	Deletes a role from the system.
<code>rolemod(1M)</code>	Modifies a role's properties on the system.
<code>smattrpop(1M)</code>	Merges the source security attribute database into the target database. For use in situations where local databases need to be merged into a name service and in upgrades where conversion scripts are not supplied.
<code>smexec(1M)</code>	Manages entries in the <code>exec_attr</code> database. Requires authentication.
<code>smmultiuser(1M)</code>	Manages bulk operations on user accounts. Requires authentication.
<code>smuser(1M)</code>	Manages user entries. Requires authentication.
<code>smprofile(1M)</code>	Manages rights profiles in the <code>prof_attr</code> and <code>exec_attr</code> databases. Requires authentication.
<code>smrole(1M)</code>	Manages roles and users in role accounts. Requires authentication.
<code>useradd(1M)</code>	Adds a user account to the system. The <code>-P</code> option assigns a role to a user's account.
<code>userdel(1M)</code>	Deletes a user's login from the system.
<code>usermod(1M)</code>	Modifies a user's account properties on the system.

## Commands That Require Authorizations

The following table provides examples of how authorizations are used to limit command options in the Solaris environment. See also "Authorizations" on page 284.

**TABLE 19-8** Commands and Associated Authorizations

Commands	Authorization Requirements
at(1)	solaris.jobs.user required for all options (when neither at.allow nor at.deny files exist)
atq(1)	solaris.jobs.admin required for all options
crontab(1)	solaris.jobs.user required for the option to submit a job (when neither crontab.allow nor crontab.deny files exist)  solaris.jobs.admin required for the options to list or modify other users' crontab files
allocate(1) (with BSM enabled only)	solaris.device.allocate (or other authorization as specified in device_allocate(4)) required to allocate a device.  solaris.device.revoke (or other authorization as specified in device_allocate file) required to allocate a device to another user (-F option)
deallocate(1) (with BSM enabled only)	solaris.device.allocate (or other authorization as specified in device_allocate(4)) required to deallocate another user's device.  solaris.device.revoke (or other authorization as specified in device_allocate) required to force deallocation of the specified device (-F option) or all devices (-I option)
list_devices(1) (with BSM enabled only)	solaris.device.revoke required to list another user's devices (-U option)

## Using the Automated Security Enhancement Tool (Tasks)

---

This chapter describes how to use the Automated Security Enhancement Tool (ASET) to monitor or restrict access to system files and directories.

This is a list of the step-by-step instructions in this chapter.

- “Automated Security Enhancement Tool (ASET)” on page 295
- “Running ASET” on page 313
- “Troubleshooting ASET Problems” on page 316

---

## Automated Security Enhancement Tool (ASET)

The Solaris 9 release includes the Automated Security Enhancement Tool (ASET). ASET helps you monitor and control system security by automatically performing tasks that you would otherwise do manually.

The ASET security package provides automated administration tools that enable you to control and monitor your system’s security. You specify a security level—low, medium, or high—at which ASET will run. At each higher level, ASET’s file-control functions increase to reduce file access and tighten your system security.

There are seven tasks involved with ASET, each task performs specific checks and adjustments to system files. The ASET tasks tighten file permissions, check the contents of critical system files for security weaknesses, and monitor crucial areas. ASET can also safeguard a network by applying the basic requirements of a firewall system to a system that serves as a gateway system. (See “Firewall Setup” on page 299.)

ASET uses master files for configuration. Master files, reports, and other ASET files are in the `/usr/aset` directory. These files can be changed to suit the particular requirements of your site.

Each task generates a report that notes detected security weaknesses and any changes the task has made to the system files. When run at the highest security level, ASET will attempt to modify all system security weaknesses. If ASET cannot correct a potential security problem, it reports the existence of the problem.

You can initiate an ASET session by using the `/usr/aset` command interactively. Or, you can set up ASET to run periodically by putting an entry into the `crontab` file.

ASET tasks are disk-intensive and can interfere with regular activities. To minimize the impact on system performance, schedule ASET to run when system activity level is lowest, for example, once every 24 or 48 hours at midnight.

## ASET Security Levels

ASET can be set to operate at one of three security levels: low, medium, or high. At each higher level, ASET's file-control functions increase to reduce file access and heighten system security. These functions range from monitoring system security without limiting users' file access, to increasingly tightening access permissions until the system is fully secured.

The following table outlines these three levels of security.

Security Level	Description
Low	Ensures that attributes of system files are set to standard release values. ASET performs several checks and reports potential security weaknesses. At this level, ASET takes no action and does not affect system services.
Medium	Provides adequate security control for most environments. ASET modifies some settings of system files and parameters, restricting system access to reduce the risks from security attacks. ASET reports security weaknesses and any modifications it makes to restrict access. At this level, ASET does not affect system services.
High	Renders a highly secure system. ASET adjusts many system files and parameter settings to minimum access permissions. Most system applications and commands continue to function normally. However, at this level, security considerations take precedence over other system behavior.



---

**Note** – ASET does not change the permissions of a file to make it less secure, unless you downgrade the security level or intentionally revert the system to the settings that existed prior to running ASET.

---

## ASET Tasks

This section discusses what ASET does. You should understand each ASET task (what its objectives are, what operations it performs, and what system components it affects) to interpret and use the reports effectively.

ASET report files contain messages that describe as specifically as possible any problems that were discovered by each ASET task. These messages can help you diagnose and correct these problems. However, successful use of ASET assumes that you possess a general understanding of system administration and system components. If you are a novice administrator, you can refer to other Solaris system administration documentation and related manual pages to prepare yourself for ASET administration.

The `taskstat` utility identifies the tasks that have been completed and the tasks that are still running. Each completed task produces a report file. For a complete description of the `taskstat` utility, refer to `taskstat(1M)`.

## System Files Permissions Tuning

This task sets the permissions on system files to the security level you designate. This task is run when the system is installed. If you decide later to alter the previously established levels, run this task again. At low security, the permissions are set to values that are appropriate for an open information-sharing environment. At medium security, the permissions are tightened to produce adequate security for most environments. At high security, they are tightened to severely restrict access.

Any modifications that this task makes to system files permissions or parameter settings are reported in the `tune.rpt` file. For example of the files that ASET consults when it sets permissions, see “Tune Files” on page 311.

## System Files Checks

This task examines system files and compares each file with a description of that file as it is listed in a master file. The master file is created the first time ASET runs this task. The master file contains the system file settings that are enforced by `checklist` for the specified security level.

A list of directories whose files are to be checked is defined for each security level. You can use the default list, or you can modify it, specifying different directories for each level.

For each file, the following criteria are checked:

- Owner and group
- Permission bits
- Size and checksum
- Number of links
- Last modification time

Any discrepancies found are reported in the `cklist.rpt` file. This file contains the results of comparing system file size, permission, and checksum values to the master file.

## User and Group Checks

This task checks the consistency and integrity of user accounts and groups as they are defined in the `passwd` and `group` files. This task checks the local, and NIS or NIS+ password files. NIS+ password file problems are reported but not corrected. This task checks for the following violations:

- Duplicate names or IDs
- Entries in incorrect format
- Accounts without a password
- Invalid login directories
- The `nobody` account
- Null group password
- A plus sign (+) in the `/etc/passwd` file on an NIS (or NIS+) server

Discrepancies are reported in the `usrgrp.rpt` file.

## System Configuration Files Check

During this task, ASET checks various system tables, most of which are in the `/etc` directory. These files are the following:

- `/etc/default/login`
- `/etc/hosts.equiv`
- `/etc/inetd.conf`
- `/etc/aliases`
- `/var/adm/utmpx`
- `/.rhosts`
- `/etc/vfstab`
- `/etc/dfs/dfstab`
- `/etc/ftpd/ftpusers`

ASET performs various checks and modifications on these files, and reports all problems in the `sysconf.rpt` file.

## Environment Variables Check

This task checks how the `PATH` and `UMASK` environment variables are set for root, and other users, in the `/.profile`, `/.login`, and `/.cshrc` files.

The results of checking the environment for security are reported in the `env.rpt` file.

## eeprom Check

This task checks the value of the `eeprom` security parameter to ensure that it is set to the appropriate security level. You can set the `eeprom` security parameter to `none`, `command`, or `full`.

ASET does not change this setting, but reports its recommendations in the `eeprom.rpt` file.

## Firewall Setup

This task ensures that the system can be safely used as a network relay. This task protects an internal network from external public networks by setting up a dedicated system as a firewall, which is described in “Firewall Systems” on page 206. The firewall system separates two networks. In this situation, each network approaches the other network as untrusted. The firewall setup task disables the forwarding of Internet Protocol (IP) packets and hides routing information from the external network.

The firewall task runs at all security levels, but takes action only at the highest level. If you want to run ASET at high security, but find that your system does not require firewall protection, you can eliminate the firewall task by editing the `asetenv` file.

Any changes that are made are reported in the `firewall.rpt` file.

## ASET Execution Log

ASET generates an execution log whether it runs interactively or in the background. By default, ASET generates the log file on standard output. The execution log confirms that ASET ran at the designated time, and also contains any execution error messages. The `aset -n` command directs the log to be delivered by electronic mail to a designated user. For a complete list of ASET options, refer to the `aset(1M)` man page.

## Example of an ASET Execution Log File

```
ASET running at security level low
```

```
Machine=example; Current time = 0325_08:00
```

```
aset: Using /usr/aset as working directory

Executing task list...
    firewall
    env
    sysconfig
    usrgroup
    tune
    cklist
    eeprom
All tasks executed. Some background tasks may still be running.

Run /usr/aset/util/taskstat to check their status:
    $/usr/aset/util/taskstat    aset_dir
Where aset_dir is ASET's operating directory, currently=/usr/aset

When the tasks complete, the reports can be found in:
    /usr/aset/reports/latest/*.rpt
You can view them by:
more /usr/aset/reports/latest/*.rpt
```

The execution log first shows the system and time that ASET was run. Then, the execution log lists each task as it is started.

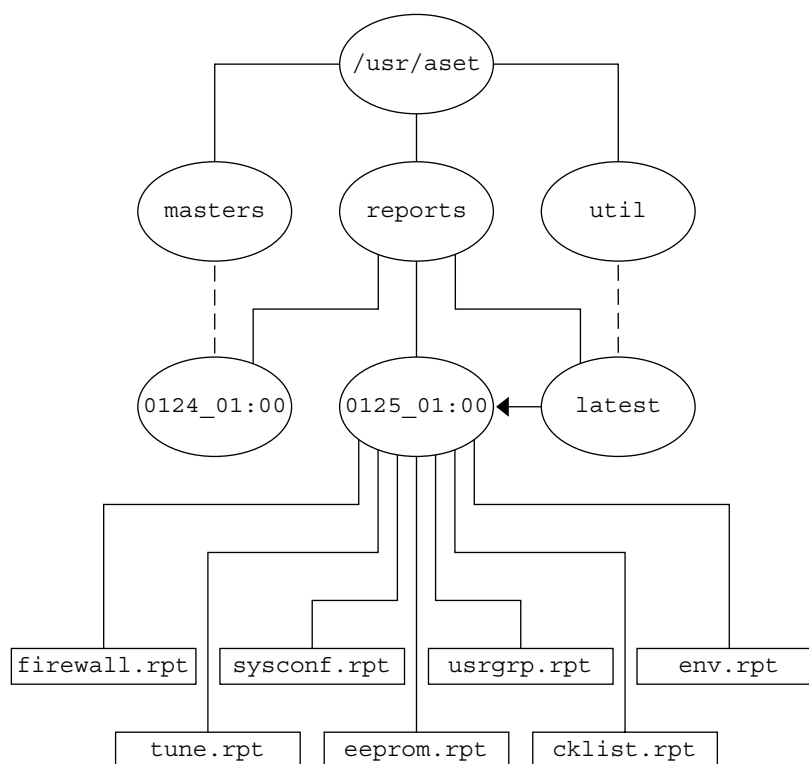
ASET invokes a background process for each of these tasks, which are described in “ASET Tasks” on page 297. The task is listed in the execution log when it starts. This listing does not indicate that it has been completed. To check the status of the background tasks, use the `taskstat` command.

## ASET Reports

All report files that are generated from ASET tasks are stored in subdirectories under the `/usr/aset/reports` directory. This section describes the structure of the `/usr/aset/reports` directory, and provides guidelines on managing the report files.

ASET places the report files in subdirectories that are named to reflect the time and date when the reports are generated. This convention enables you to keep an orderly trail of records that document the system status as it varies between ASET executions. You can monitor and compare these reports to determine the soundness of your system's security.

The following figure shows an example of the `reports` directory structure.



**FIGURE 20-1** Structure of the ASET reports Directory

This example shows two report subdirectories.

- 0124\_01:00
- 0125\_01:00

The subdirectory names indicate the date and time that the reports were generated. Each report subdirectory name has the following format:

*monthdate\_hour:minute*

where *month*, *date*, *hour*, and *minute* are all two-digit numbers. For example, 0125\_01:00 represents January 25, at 1 a.m.

Each of the two report subdirectories contains a collection of reports that are generated from one execution of ASET.

The `latest` directory is a symbolic link that always points to the subdirectory that contains the latest reports. Therefore, to look at the latest reports that ASET has generated, you can go to the `/usr/aset/reports/latest` directory. There is a report file in this directory for each task that ASET performed during its most recent execution.

## Format of ASET Report Files

Each report file is named after the task that generates it. See the following table for a list of tasks and their reports.

**TABLE 20-1** ASET Tasks and Resulting Reports

Tasks	Report
System files permissions tuning ( <code>tune</code> )	<code>tune.rpt</code>
System files checks ( <code>cklist</code> )	<code>cklist.rpt</code>
User and group checks ( <code>usrgrp</code> )	<code>usrgrp.rpt</code>
System configuration files check ( <code>sysconf</code> )	<code>sysconf.rpt</code>
Environment variables check ( <code>env</code> )	<code>env.rpt</code>
EEPROM check ( <code>eeeprom</code> )	<code>eeeprom.rpt</code>
Firewall setup ( <code>firewall</code> )	<code>firewall.rpt</code>

Within each report file, messages are bracketed by a beginning and an ending banner line. Sometimes, a task terminates prematurely; for example, when a component of ASET is accidentally removed or damaged. In such cases, the report file usually contains a message near the end that indicates the reason for the premature termination.

The following is a sample report file, `usrgrp.rpt`.

```
*** Begin User and Group Checking ***

Checking /etc/passwd ...
Warning! Password file, line 10, no passwd
:sync::1:1:::/bin/sync
..end user check; starting group check ...
Checking /etc/group...
*** End User And group Checking ***
```

## Examining ASET Report Files

After you initially run or reconfigure ASET, you should examine the report files closely. Reconfiguration includes modifying the `aset.env` file or the master files in the `masters` subdirectory, or changing the security level at which ASET operates.

The reports record any errors that were introduced when you reconfigured ASET. By watching the reports closely, you can react to, and solve, problems as they arise.

## Comparing ASET Report Files

After you monitor the report files for a period during which there are no configuration changes or system updates, you might find that the content of the reports begin to stabilize and that it contains little, if any, unexpected information. You can use the `diff` utility to compare reports.

## ASET Master Files

ASET's master files, `tune.high`, `tune.low`, `tune.med`, and `uid_aliases`, are located in the `/usr/aset/masters` directory. ASET uses the master files to define security levels.

## Tune Files

The `tune.low`, `tune.med`, and `tune.high` master files define the available ASET security levels. They specify the attributes of system files at each level and are used for comparison and reference purposes.

## The `uid_aliases` File

The `uid_aliases` file contains a list of multiple user accounts that share the same user ID (UID). Normally, ASET warns about such multiple user accounts because this practice lessens accountability. You can allow for exceptions to this rule by listing the exceptions in the `uid_aliases` file. ASET does not report entries in the `passwd` file with duplicate UIDs if these entries are specified in the `uid_aliases` file.

Avoid having multiple user accounts (password entries) share the same UID. You should consider other methods of achieving your objective. For example, if you intend for several users to share a set of permissions, you could create a group account. The sharing of UIDs should be your last resort, used only when absolutely necessary and when other methods will not accomplish your objectives.

You can use the `UID_ALIASES` environment variable to specify an alternate aliases file. The default file is `/usr/aset/masters/uid_aliases`.

## The Checklist Files

The master files that are used by the systems files checks are generated when you first execute ASET, or when you run ASET after you change the security level.

The following environment variables define the files that are checked by this task:

- CKLISTPATH\_LOW
- CKLISTPATH\_MED
- CKLISTPATH\_HIGH

## ASET Environment File (`asetenv`)

The environment file, `asetenv`, contains a list of environment variables that affect ASET tasks. Some of these variables can be changed to modify ASET operation.

## Configuring ASET

This section discusses how ASET is configured and the environment under which it operates.

ASET requires minimum administration and configuration, and in most cases, you can run it with the default values. You can, however, fine-tune some of the parameters that affect the operation and behavior of ASET to maximize its benefit. Before you change the default values, you should understand how ASET works, and how it affects the components of your system.

ASET relies on four configuration files to control the behavior of its tasks:

- `/usr/aset/asetenv`
- `/usr/aset/masters/tune.low`
- `/usr/aset/masters/tune.med`
- `/usr/aset/masters/tune.high`

## Modifying the Environment File (`asetenv`)

The `/usr/aset/asetenv` file has two main sections:

- A user-configurable environment variables section
- An internal environment variables section

You can alter the user-configurable parameters section. However, the settings in the internal environment variables section are for internal use only and should not be modified.

You can edit the entries in the user-configurable section to do the following:

- Choose which tasks to run
- Specify the directories for the system files checks task
- Schedule ASET execution
- Specify a UID aliases file
- Extend checks to NIS+ tables



## Choose Which Tasks to Run: TASKS

Each task that ASET performs monitors a particular area of system security. In most system environments, all the tasks are necessary to provide balanced security coverage. However, you might decide to eliminate one or more tasks.

For example, the firewall task runs at all security levels, but takes action only at the high security level. You might want to run ASET at the high security level, but you do not require firewall protection.

You can set up ASET to run at the high security level without the firewall feature by editing the `TASKS` list of environment variables in the `asetenv` file. By default, the `TASKS` list contains all of the ASET tasks. To delete a task, remove the task-related environment variable from the file. In this case, you would delete the `firewall` environment variable from the list. The next time ASET runs, the excluded task will not be performed.

The following example shows the `TASKS` list with all of the ASET tasks.

```
TASKS="env sysconfig usrgrp tune cklist eeprom firewall"
```

## Specify Directories for System Files Checks Task: CKLISTPATH

The system files check checks the attributes of files in selected system directories. You define which directories to check by using these checklist path environment variables.

The `CKLISTPATH_LOW` variable defines the directories to be checked at the low security level. `CKLISTPATH_MED` and `CKLISTPATH_HIGH` environment variables function similarly for the medium and high security levels.

The directory list that is defined by an environment variable at a lower security level should be a subset of the directory list that is defined at the next higher level. For example, all directories that are specified for `CKLISTPATH_LOW` should be included in `CKLISTPATH_MED`, and all the directories that are specified for `CKLISTPATH_MED` should be included in `CKLISTPATH_HIGH`.

Checks performed on these directories are not recursive. ASET only checks those directories that are explicitly listed in the environment variable. ASET does not check their subdirectories.

You can edit these environment variable definitions to add or delete directories that you want ASET to check. Note that these checklists are useful only for system files that do not normally change from day to day. A user's home directory, for example, is generally too dynamic to be a candidate for a checklist.

## Schedule ASET Execution: PERIODIC\_SCHEDULE

When you start ASET, you can start it interactively, or use the `-p` option to request that the ASET tasks run at a scheduled time. You can run ASET periodically, at a time when system demand is light. For example, ASET consults `PERIODIC_SCHEDULE` to determine how frequently to execute the ASET tasks, and at what time to run them. For detailed instructions about setting up ASET to run periodically, see “How to Run ASET Periodically” on page 314.

The format of `PERIODIC_SCHEDULE` follows the format of `crontab` entries. See `crontab(1)` for complete information.

## Specify an Aliases File: UID\_ALIASES

The `UID_ALIASES` variable specifies an aliases file that lists shared UIDs. The default file is `/usr/aset/masters/uid_aliases`.

## Extend Checks to NIS+ Tables: YPCHECK

The `YPCHECK` environment variable specifies whether ASET should also check system configuration file tables. `YPCHECK` is a Boolean variable; you can specify only `true` or `false` for it. The default value is `false`, which disables NIS+ table checking.

To understand how this environment variable works, consider its effect on the `passwd` file. When set to `false`, ASET checks the local `passwd` file. When set to `true`, the task also checks the NIS+ `passwd` file for the domain of the system.

---

**Note** – Although ASET automatically repairs the local tables, it only reports potential problems in the NIS+ tables. ASET does not change them.

---

## Modifying the Tune Files

ASET uses the three master tune files, `tune.low`, `tune.med`, and `tune.high`, to ease or tighten access to critical system files. These master files are located in the `/usr/aset/masters` directory, and they can be modified to suit your environment. For additional information, see “Tune Files” on page 311.

The `tune.low` file sets permissions to values that are appropriate for default system settings. The `tune.med` file further restricts these permissions and includes entries that are not present in `tune.low`. The `tune.high` file restricts permissions even further.

---

**Note** – Modify settings in the tune files by adding or deleting file entries. Setting a permission to a less restrictive value than the current setting has no effect. The ASET tasks do not relax permissions unless you downgrade your system security to a lower level.

---

## Restoring System Files Modified by ASET

When ASET is executed for the first time, it saves and archives the original system files. The `aset.restore` utility reinstates these files. This utility also deschedules ASET, if it is currently scheduled for periodic execution. The `aset.restore` command is located in `/usr/aset`, the ASET operating directory.

Changes that are made to system files are lost when you run the `aset.restore` command.

You should use the `aset.restore` command in the following instances:

- When you want to remove ASET changes and restore the original system. If you want to deactivate ASET permanently, you can remove it from `cron` scheduling if the `aset` command had previously been added to root's `crontab`. For instructions on how to use `cron` to remove automatic execution, see “How to Stop Running ASET Periodically” on page 315.
- After a brief period of experimenting with ASET, to restore the original system state.
- When some major system feature is not working properly, and you suspect that ASET is causing the problem.

## Network Operation With the NFS System

Generally, ASET is used in standalone mode, even on a system that is part of a network. As system administrator for your standalone system, you are responsible for the security of your system and for running and managing ASET to protect your system.

You can also use ASET in the NFS distributed environment. As a network administrator, you are responsible for installing, running, and managing various administrative tasks for all your clients. To facilitate ASET management across several client systems, you can make configuration changes that are applied globally to all clients, which eliminates the need for you to log in to each system to repeat the process.

When you are deciding how to set up ASET on your networked systems, you should consider how much you want users to control security on their own systems, and how much you want to centralize responsibility for security control.

## Providing a Global Configuration for Each Security Level

A situation might arise where you want to set up more than one network configuration. For example, you might want to set up one configuration for clients that are designated with low security level, another configuration for those clients with medium level, and yet another configuration with high level.

If you need to create a separate ASET network configuration for each security level, you can create three ASET configurations on the server, one configuration for each level. You would export each configuration to the clients with the appropriate security level. Some ASET components that are common to all three configurations could be shared by using links.

## Collecting ASET Reports

Not only can you centralize the ASET components on a server to be accessed by clients with or without superuser privilege, but you can also set up a central directory on a server to collect all reports that are produced by tasks that run on various clients. For instructions on setting up a collection mechanism, see “How to Collect ASET Reports on a Server” on page 315.

Setting up the collection of reports on a server allows you to review reports for all clients from one location. You can use this method whether or not a client has superuser privilege. Alternatively, you can leave the reports directory on the local system when you want users to monitor their own ASET reports.

## ASET Environment Variables

The following table lists the ASET environment variables and the values that they specify.

**TABLE 20-2** ASET Environment Variables and Their Meanings

Environment Variable	Value Specified
ASETDIR	ASET working directory
ASETSECLEVEL	Security level
PERIODIC_SCHEDULE	Periodic schedule
TASKS	Tasks to run
UID_ALIASES	Aliases file
YPCHECK	Whether to extend checks to NIS maps and NIS+ tables

**TABLE 20-2** ASET Environment Variables and Their Meanings (Continued)

Environment Variable	Value Specified
CKLISTPATH_LOW	Directory lists for low security
CKLISTPATH_MED	Directory list for medium security
CKLISTPATH_HIGH	Directory list for high security

The environment variables that are listed in the following sections are found in the `/usr/aset/asetenv` file. The `ASETDIR` and `ASETSECLEVEL` variables are optional and can be set only through the shell by using the `aset` command. The other environment variables can be set by editing the file.

## ASETDIR Environment Variable

`ASETDIR` specifies an ASET working directory.

From the C shell, type:

```
% setenv ASETDIR pathname
```

From the Bourne shell or the Korn shell, type:

```
$ ASETDIR=pathname  
$ export ASETDIR
```

Set *pathname* to the full path name of the ASET working directory.

## ASETSECLEVEL Environment Variable

The `ASETSECLEVEL` variable specifies a security level at which ASET tasks are executed.

From the C shell, type:

```
% setenv ASETSECLEVEL level
```

From the Bourne shell or the Korn shell, type:

```
$ ASETSECLEVEL=level  
export ASETSECLEVEL
```

In these commands, *level* can be set to one of the following:

low	Low security level
med	Medium security level

high

High security level

## PERIODIC\_SCHEDULE Environment Variable

The value of `PERIODIC_SCHEDULE` follows the same format as the `crontab` file. Specify the variable value as a string of five fields enclosed in double quotation marks, with each field separated by a space:

*"minutes hours day-of-month month day-of-week"*

**TABLE 20-3** Periodic\_Schedule Variable Values

Variable	Value
<i>minutes hours</i>	Specifies start time in number of minutes (0–59) after the hour and the hour (0–23)
<i>day-of-month</i>	Specifies the day of the month when ASET should be run, with values from 1–31
<i>month</i>	Specifies the month of the year when ASET should be run, with values from 1–12
<i>day-of-week</i>	Specifies the day of the week when ASET should be run, with values from 0–6; Sunday is day 0

The following rules apply:

- You can specify a list of values, each delimited by a comma, for any field.
- You can specify a value as a number, or you can specify it as a range; that is, a pair of numbers that are joined by a hyphen. A range states that the ASET tasks should be executed for every time that is included in the range.
- You can specify an asterisk (\*) as the value of any field. An asterisk inclusively specifies all possible values of the field.

The default entry for the `PERIODIC_SCHEDULE` variable causes ASET to execute at 12:00 midnight every day:

```
PERIODIC_SCHEDULE="0 0 * * *"
```

## TASKS Environment Variable

The `TASKS` variable lists the tasks that ASET performs. The default is to list all seven tasks:

```
TASKS="env sysconfig usrgroup tune cklist eeprom firewall"
```

## UID\_ALIASES Environment Variable

The `UID_ALIASES` variable specifies an aliases file. If present, ASET consults this file for a list of permitted multiple aliases. The format is `UID_ALIASES=pathname`, where *pathname* is the full path name of the aliases file.

The default is as follows:

```
UID_ALIASES=${ASETDIR}/masters/uid_aliases
```

## YPCHECK Environment Variable

The `YPCHECK` variable extends the task of checking system tables to include NIS or NIS+ tables. This variable is a Boolean variable, which can be set to either true or false.

The default is false, which confines the checking to local system tables:

```
YPCHECK=false
```

## CKLISTPATH\_level Environment Variable

The three checklist path variables list the directories to be checked by the system files checks task. The following definitions of the variables are set by default. They illustrate the relationship between the variables at different levels:

```
CKLISTPATH_LOW=${ASETDIR}/tasks:${ASETDIR}/util:${ASETDIR}/masters:  
/etc  
CKLISTPATH_MED=${CKLISTPATH_LOW}:/usr/bin:/usr/ucb  
CKLISTPATH_HIGH=${CKLISTPATH_MED}:/usr/lib:/sbin:/usr/sbin:/usr/ucblib
```

The values for the checklist path environment variables are similar to those values of the shell path variables, in that they are lists of directory names that are separated by colons. You use an equal sign (=) to connect the variable name to its value.

## ASET File Examples

This section has examples of some ASET files, including the tune files and the aliases file.

### Tune Files

ASET maintains three tune files. The following table describes the format of entries in all three tune files.

**TABLE 20-4** Entry Format for Tune Files

Field Name	Description
<i>pathname</i>	The full path name of the file
<i>mode</i>	A five-digit number that represents the permission setting
<i>owner</i>	The owner of the file
<i>group</i>	The group owner of the file
<i>type</i>	The type of file

The following rules apply when you edit the tune files:

- You can use regular shell wildcard characters, such as an asterisk (\*) and a question mark (?), in the path name for multiple references. See `sh(1)` for more information.
- *mode* represents the least restrictive value. If the current setting is already more restrictive than the specified value, ASET does not loosen the permission settings. For example, if the specified value is `00777`, the permission remains unchanged, because `00777` is always less restrictive than whatever the current setting is.
 

This process is how ASET handles mode setting, unless the security level is being downgraded or you are removing ASET. When you decrease the security level from what it was for the previous execution, or when you want to restore the system files to the state they were in before ASET was first executed, ASET recognizes what you are doing and decreases the protection level.
- You must use names for *owner* and *group* instead of numeric IDs.
- You can use a question mark (?) in place of *owner*, *group*, and *type* to prevent ASET from changing the existing values of these parameters.
- *type* can be `symlink` (symbolic link), `directory`, or `file` (everything else).
- Higher security level tune files reset file permissions to be at least as restrictive as they are at lower levels. Also, at higher security levels, additional files are added to the list.
- A file can match more than one tune file entry. For example, `etc/passwd` matches the `etc/pass*` and `/etc/*` entries.
- Where two entries have different permissions, the file permission is set to the most restrictive value. In the following example, the permission of the `/etc/passwd` file will be set to `00755`, which is the more restrictive of `00755` and `00770`.
 

```
/etc/pass* 00755 ? ? file
/etc/* 00770 ? ? file
```
- If two entries have different *owner* or *group* designations, the last entry takes precedence. In the following example, the owner of `/usr/sbin/chroot` will be set to `root`.



```
/usr/sbin/chroot 00555 bin bin file
/usr/sbin/chroot 00555 root bin file
```

## Aliases File

The aliases file contains a list of aliases that share the same user ID.

Each entry is in this form:

```
uid=alias1=alias2=alias3= . . .
```

<i>uid</i>	Shared UID.
<i>aliasn</i>	User account that share the UID.

For example, the following entry lists the UID 0 that is being shared by the `sysadm` and `root` accounts:

```
0=root=sysadm
```

---

## Running ASET

This section describes how to run ASET either interactively or periodically.

### ▼ How to Run ASET Interactively

1. **Become superuser.**
2. **Run ASET interactively by using the `aset` command.**

```
# /usr/aset/aset -l level -d pathname
```

<i>level</i>	Specifies the level of security. Valid values are <code>low</code> , <code>medium</code> , or <code>high</code> . The default setting is <code>low</code> . For detailed information about security levels see “ASET Security Levels” on page 296.
<i>pathname</i>	Specifies the working directory for ASET. The default is <code>/usr/aset</code> .

3. **Verify that ASET is running by viewing the ASET execution log that is displayed on the screen.**

The execution log message identifies which tasks are being run.

## Example—Running ASET Interactively

The following example shows ASET being run at low security with the default working directory.

```
# /usr/aset/aset -l low
===== ASET Execution Log =====

ASET running at security level low

Machine = jupiter; Current time = 0111_09:26

aset: Using /usr/aset as working directory

Executing task list ...
    firewall
    env
    sysconf
    usrgrp
    tune
    cklist
    eeprom

All tasks executed. Some background tasks may still be running.

Run /usr/aset/util/taskstat to check their status:
/usr/aset/util/taskstat [aset_dir]

where aset_dir is ASET's operating
directory, currently=/usr/aset.

When the tasks complete, the reports can be found in:
/usr/aset/reports/latest/*.rpt

You can view them by:
more /usr/aset/reports/latest/*.rpt
```

## ▼ How to Run ASET Periodically

### 1. Become superuser.

### 2. If necessary, set up the time when you want ASET to run periodically.

You should have ASET run when system demand is light. The `PERIODIC_SCHEDULE` environment variable in the `/usr/aset/asetenv` file is used to set up the time for ASET to run periodically. By default, the time is set for every day at midnight.

If you want to set up a different time, edit the `PERIODIC_SCHEDULE` variable in the `/usr/aset/asetenv` file. For detailed information about setting the `PERIODIC_SCHEDULE` variable see “`PERIODIC_SCHEDULE` Environment Variable”

on page 310.

**3. Add an entry to the crontab file by using the `aset` command.**

```
# /usr/aset/aset -p
```

The `-p` option inserts a line in the `crontab` file that starts ASET running at the time determined by the `PERIODIC_SCHEDULE` environment variable in the `/usr/aset/asetenv` file.

**4. Display the crontab entry to verify when ASET will run.**

```
# crontab -l root
```

## ▼ How to Stop Running ASET Periodically

**1. Become superuser.**

**2. Edit the `crontab` file.**

```
# crontab -e root
```

**3. Delete the ASET entry.**

**4. Save the changes and exit.**

**5. Display the `crontab` entry to verify that the ASET entry is deleted.**

```
# crontab -l root
```

## ▼ How to Collect ASET Reports on a Server

**1. Become superuser.**

**2. Set up a directory on the server:**

**a. Change to the `/usr/aset` directory.**

```
mars# cd /usr/aset
```

**b. Create a `rptdir` directory.**

```
mars# mkdir rptdir
```

**c. Change to the `rptdir` directory, and create a `client_rpt` directory.**

This creates a subdirectory (`client_rpt`) for a client. Repeat this step for each client whose reports you need to collect.

```
mars# cd rptdir
```

```
mars# mkdir client_rpt
```

The following example shows the creation of the directory `all_reports`, and the subdirectories `pluto_rpt` and `neptune_rpt`.

```
mars# cd /usr/aset
mars# mkdir all_reports
mars# cd all_reports
mars# mkdir pluto_rpt
mars# mkdir neptune_rpt
```

**3. Add the `client_rpt` directories to the `/etc/dfs/dfstab` file.**

The directories should have read and write options.

For example, the following entries in the `dfstab` file are shared with read and write permissions.

```
share -F nfs -o rw=pluto /usr/aset/all_reports/pluto_rpt
share -F nfs -o rw=neptune /usr/aset/all_reports/neptune_rpt
```

**4. Make the resources in the `dfstab` file available to the clients.**

```
# shareall
```

**5. On each client, mount the client subdirectory from the server at the mount point, `/usr/aset/masters/reports`.**

```
# mount server:/usr/aset/client_rpt /usr/aset/masters/reports
```

**6. Edit the `/etc/vfstab` file to mount the directory automatically at boot time.**

The following sample entry in `/etc/vfstab` on `neptune` lists the directory to be mounted from `mars`, `/usr/aset/all_reports/neptune_rpt`, and the mount point on `neptune`, `/usr/aset/reports`. At boot time, the directories that are listed in `vfstab` are automatically mounted.

```
mars:/usr/aset/all_reports/neptune.rpt /usr/aset/reports nfs - yes hard
```

---

## Troubleshooting ASET Problems

This section documents the error messages that are generated by ASET.

### ASET Error Messages

ASET failed: no mail program found.

**Cause:** ASET is directed to send the execution log to a user, but no mail program can be found.

**Solution:** Install a mail program.

Usage: `aset [-n user[@host]] in /bin/mail or /usr/ucb/mail.`

Cannot decide current and previous security levels.

**Cause:** ASET cannot determine what the security levels are for the current and previous invocations.

**Solution:** Ensure the current security level is set either through the command-line option or the `ASETSECLEVEL` environment variable. Also, ensure that the last line of `ASETDIR/archives/asetseclevel.arch` correctly reflects the previous security level. If these values are not set or are incorrect, correct them.

ASET working directory undefined.

To specify, set `ASETDIR` environment variable or use command line option `-d`.

ASET startup unsuccessful.

**Cause:** The ASET working (operating) directory is not defined, or it is defined incorrectly.

**Solution:** Use the `ASETDIR` environment variable or the `-d` command-line option to correct it, and restart ASET.

ASET working directory `$ASETDIR` missing.

ASET startup unsuccessful.

**Cause:** The ASET working (operating) directory is not defined, or it is defined incorrectly. This problem might be because the `ASETDIR` variable or the `-d` command-line option refers to a nonexistent directory.

**Solution:** Ensure that the correct directory, that is, the directory that contains the ASET directory hierarchy, is referred to correctly.

Cannot expand `$ASETDIR` to full pathname.

**Cause:** ASET cannot expand the directory name that is given by the `ASETDIR` variable or the `-d` command-line option to a full path name.

**Solution:** Ensure that the directory name is correct, and that it refers to an existing directory to which the user has access.

`aset: invalid/undefined security level.`

To specify, set `ASETSECLEVEL` environment variable or use command line option `-l`, with argument= `low/med/high`.

**Cause:** The security level is not defined, or it is defined incorrectly. Only the values `low`, `med`, or `high` are acceptable.

**Solution:** Use the `ASETSECLEVEL` variable or the `-l` command-line option to specify one of the three values.

ASET environment file `asetenv` not found in `$ASETDIR`.

ASET startup unsuccessful.

**Cause:** ASET cannot locate an `asetenv` file in its working directory.

**Solution:** Ensure there is an `asetenv` file in ASET's working directory. See `asetenv(4)` for the details about this file.

`filename` doesn't exist or is not readable.

**Cause:** The file that is referred to by `filename` doesn't exist or is not readable. This problem can occur when you are using the `-u` option, where you can specify a file that contains a list of users whom you want to check.

**Solution:** Ensure that the argument to the `-u` option exists and is readable.

ASET task list `TASKLIST` undefined.

**Cause:** The ASET task list, which should be defined in the `asetenv` file, is not defined. This message can mean that your `asetenv` file is bad.

**Solution:** Examine your `asetenv` file. Ensure that the task list is defined in the `User Configurable` section. Also check other parts of the file to ensure that the file is intact. See `asetenv(4)` for the content of a valid `asetenv` file.

ASET task list `$TASKLIST` missing.

ASET startup unsuccessful.

**Cause:** The ASET task list, which should be defined in the `asetenv` file, is not defined. This message can mean that your `asetenv` file is bad.

**Solution:** Examine your `asetenv` file. Ensure that the task list is defined in the `User Configurable` section. Also check other parts of the file to ensure that the file is intact. See `asetenv(4)` for the content of a valid `asetenv` file.

`Schedule` undefined for periodic invocation.

No tasks executed or scheduled. Check `asetenv` file.

**Cause:** ASET scheduling is requested by using the `-p` option, but the environment variable `PERIODIC_SCHEDULE` is undefined in the `asetenv` file.

**Solution:** Check the `User Configurable` section of the `asetenv` file to ensure that the variable is defined and is in proper format.

Warning! Duplicate ASET execution scheduled.

Check `crontab` file.

**Cause:** ASET is scheduled to run more than once. In other words, ASET scheduling is requested while a schedule is already in effect. This message does not necessarily indicate an error if more than one schedule is indeed desired. In this instance, the messages servers only as a warning. If you want more than one schedule, you should use the proper scheduling format with the `crontab` command. For more information, see the `crontab(1)` man page.

**Solution:** Verify, through the `crontab` command, that the correct schedule is in effect. Ensure that no unnecessary `crontab` entries for ASET are in place.





## Auditing Topics

---

Chapter 22	Provides overview information about auditing.
Chapter 23	Provides information to help with auditing planning issues.
Chapter 24	Provides step-by-step instructions to configure and manage auditing.
Chapter 25	Provides information about the files associated with auditing and the structure of the audit tokens.



## BSM (Overview)

---

The Basic Security Module (BSM) provides two security features. The first feature is a security auditing system, which includes tools to assist with the analysis of the auditing data. The second feature is a device-allocation mechanism, which provides the required object-reuse characteristics for removable or assignable devices.

This chapter introduces the concepts behind BSM. This is a list of the overview information in this chapter.

- “What Is Auditing?” on page 323
- “How Is Auditing Related to Security?” on page 325
- “BSM Terminology” on page 325
- “How Does Auditing Work?” on page 324

---

## What Is Auditing?

Auditing is the collection of data that provides a record of security-related system events. This data can then be used to assign responsibility to actions that take place on a host. Successful auditing starts with two security features: identification and authentication. At login, after a user supplies a user name and password, a unique audit ID is associated with the user’s process. The audit ID is inherited by every process that is started during the login session. Even if a user changes identity (see the `su(1M)` man page), all user actions are tracked with the same audit ID.

Auditing makes it possible to do the following:

- Monitor security-relevant events that take place on the host
- Record the events in a network-wide audit trail
- Detect misuse or unauthorized activity
- Review patterns of access, and see the access histories of individuals and objects

- Discover attempts to bypass the protection mechanisms
- Discover extended use of privilege that occurs when a user changes identity

During system configuration, you select which activities to monitor. You can also fine-tune the degree of auditing that is done for individual users.

After audit data is collected, audit-reduction and interpretation tools allow you to examine interesting parts of the audit trail. For example, you can choose to review audit records for individual users or groups, to examine all records for a certain type of event on a specific day, or to select records that were generated at a certain time of day.

---

## How Does Auditing Work?

Auditing is the generation of audit records when specified events occur. Most commonly, events that generate audit records include the following:

- System startup and shutdown
- Login and logout
- Process creation or destruction, or thread creation or destruction
- Opening, closing, creating, destroying, or renaming of objects
- Use of privilege, identification, and authentication actions
- Discretionary Access Control (DAC) changes by a process or user
- Installation-specific administrative actions

Audit records are generated from three sources:

- By an application
- As a result of an asynchronous event
- As a result of a process system call

Once the relevant event information has been selected, the information is formatted into an audit record. The record is then placed in a kernel buffer known as the audit queue. From this temporary location within the kernel, audit records are written to audit files. Where the audit files are located is determined by entries in the `audit_control` file. The location can include multiple partitions on the same machine, partitions on different machines, or partitions on machines on different but linked networks. The collection of audit files that are linked together is considered an audit trail.

Audit records accumulate in audit files chronologically. Contained in each audit record is information that identifies the event, what caused the event, the time of the event, and other relevant information.

---

## How Is Auditing Related to Security?

Securing a computer system, especially a system on a network, involves mechanisms that control activities before system or user processes begin, that monitor activities as they occur, and that report activities after they have happened. While the setting up of auditing requires that parameters be set before users log in or machine processes begin, most auditing activities involve monitoring current events and reporting those events that meet the specified parameters. How auditing monitors and reports these events is discussed in detail in Chapter 23 and Chapter 24.

Auditing cannot prevent hackers from unauthorized entry. However, auditing can report, for example, that a specific user performed specific actions at a specific time and date, and can identify the user by entry path and user name. Such information can be reported immediately to your terminal and to a file for later analysis. Thus, auditing provides data that helps you determine both how system security was compromised and what loopholes need to be closed to ensure the desired level of security.

---

## BSM Terminology

The following terms are used to describe the BSM service. Some definitions include pointers to more complete descriptions.

**TABLE 22-1** BSM Terms

Term	Definition
Audit class	A grouping of audit events. Audit classes provide a way to manage a group of events. For more information, see "Audit Classes" on page 327.
Audit directory	A repository of audit files. For a description of the types of audit directories, see "Audit Directory" on page 328.
Audit event	A security-related system action that is audited. For a discussion of the types of audit events, see "Audit Events" on page 326.
Audit flag	A variable that is used to determine which classes of events to audit and when to audit them. For more information about audit flags, see "Audit Flags" on page 327.

**TABLE 22-1** BSM Terms (Continued)

Term	Definition
Audit policy	A set of auditing options that you can enable or disable for a particular configuration. These options include whether to record certain kinds of audit data or whether to suspend auditable actions when the audit trail is full.
Audit record	The data, that is stored in binary form, that describes a single audit event. An audit record is composed of audit tokens. For more information about audit records, see "Audit Records" on page 328.
Audit token	A set of data, that is stored in binary form, that describes an attribute of an audit event, such as a process, a path, or other object. For descriptions of all the audit tokens, see "Audit Token Formats" on page 374.
Audit trail	A collection of one or more audit files that might reside in separate audit file partitions.
Device allocation	A mechanism that enables you to restrict use of a device, as well as to erase any leftover data after a device has been used. For a description of device allocation, see "Device Allocation" on page 328.

## Audit Events

Security-relevant system actions can be audited. These auditable actions are defined as *audit events*. Audit events are listed in the `/etc/security/audit_event` file. Each auditable event is defined in the file by a symbolic name, an event number, a set of preselection classes, and a short description (see the `audit_event(4)` man page).

There are several categories of audit events. The primary distinction is between events that are generated by the kernel (called kernel-level events), and events that are generated by applications (called user-level events). Whether the event is generated by the kernel or by a user-level application determines the number range of the event number that identifies the event.

**TABLE 22-2** Audit Event Categories

Number Range	Type of Event
1–2047	Kernel-level audit events
2048–65535	User-level audit events
2048–32767	Reserved for SunOS user-level programs
32768–65535	Available for third-party applications

## Kernel-Level Audit Events

Events that are generated by the kernel (system calls) have event numbers between 1 and 2047. The event names for kernel events begin with `AUE_`, followed by an uppercase mnemonic for the event. For example, the event number for the `creat()` system call is 4, and the event name is `AUE_CREAT`.

## User-Level Audit Events

Events that are generated by application software (outside the kernel) range from 2048 to 65535. The event names begin with `AUE_`, followed by a lowercase mnemonic for the event. For example, the event number for the `rlogin` command is 6155, and the event name is `AUE_rlogin`. Table 22-2 shows general categories of user-related events.

## Nonattributable Audit Events

Most events are attributable to an individual user, but not all. These events are known as *nonattributable events*. Events are nonattributable if they occur at the kernel-interrupt level or before a user is identified and authenticated. Nonattributable events are auditable as well. `AUE_ENTERPROM` (kernel-level event number 153), and `AUE_mountd_mount` (user-level event number 6156), are examples of nonattributable events. Check the `/etc/security/audit_event` file for the exact numbers of individual events.

## Audit Classes

Each audit event is also defined as belonging to an *audit class* or classes. By assigning events to classes, you can more easily deal with large numbers of events. When you name a class, you simultaneously address all the events in that class. The mapping of audit events to classes is configurable. These configuration changes can be made in the `audit_event` file.

Whether an auditable event is recorded in the audit trail depends on whether you preselect a class for auditing that includes the specific event. There are 32 possible audit classes. The classes include the two global classes: `all` and `no`. The audit classes are described in Table 25-2.

## Audit Flags

*Audit flags* indicate classes of events to audit. Machine-wide defaults for auditing are specified for all users on each machine by flags in the `audit_control` file, which is described in “The `audit_control` File” on page 364.

You can modify what is audited for individual users by putting audit flags in a user's entry in the `audit_user` file. The audit flags are also used as arguments to the `auditconfig`. See the `auditconfig(1M)` man page.

## Audit Records

Each *audit record* describes the occurrence of a single audited event and includes information such as who did the action, which files were affected, what action was attempted, and where and when the action occurred.

The type of information that is saved for each audit event is defined as a set of *audit tokens*. Each time an audit record is created for an event, the record contains some or all of the tokens that are defined for it, depending on the nature of the event. You can generate audit record descriptions with the `auditlist` command. For more information, see "How to Display Audit Record Formats" on page 351. For a description of the structure of each audit token, see "Audit Token Formats" on page 374.

Audit records are collected in an audit trail (see the `audit.log(4)` man page) and can be converted to a readable format by the `praudit` command (see the `praudit(1M)` man page). See "The `praudit` Command" on page 360 for details.

## Audit Directory

An *audit directory* holds a collection of audit files. Many audit directories are used in a typical installation. The three types of audit directories are as follows:

- **Primary audit directory** – A directory where the audit files for a system are placed under normal conditions.
- **Secondary audit directory** – A directory where the audit files for a system are placed if the primary directory is full or not available.
- **Directory of last resort** – A local audit directory that is used if both the primary audit directory and all secondary audit directories are not available.

## Device Allocation

The device-allocation mechanism enables you to restrict use of a device, as well as to erase any leftover data after a device has been used. These actions increase the security of a device. For more information see "Managing Device Allocation" on page 352 or "Device Allocation Reference" on page 392.



---

## Audit Planning

---

This chapter describes how to set up auditing for your Solaris installation. In particular, the chapter covers issues that you need to consider before you enable the auditing service. This chapter covers the following:

- “Handling the Audit Trail” on page 329
- “Deciding Who and What to Audit” on page 330
- “Controlling Auditing Costs” on page 333
- “Auditing Efficiently” on page 334

---

## Handling the Audit Trail

File space for the audit trail is one of the biggest issues with auditing. Each host should have several audit directories that are configured for audit files. You should decide how to configure the audit directories as one of the first steps before you enable auditing on any hosts. The following table lists the issues that should be resolved when you plan for audit trail storage.

Issue	Description
1. Determine how much auditing your site needs	Balance your site’s security needs against the availability of disk space for the audit trail. For guidance on how to reduce space requirements while still maintaining site security, as well as how to design audit storage, see “Controlling Auditing Costs” on page 333 and “Auditing Efficiently” on page 334.

Issue	Description
2. Determine which systems are to be audited and which systems are to store audit files	Decide which hosts in your network need to be audited. Make sure to create at least one local audit directory for each host that is to be audited. Then, decide which hosts are to hold most of the audit trail.
3. Determine the names and locations of the audit directories	Create a list of all the audit directories that you plan to use.
4. Plan which hosts are to use which audit directories	Create a map that shows which host should use which audit directory. This step helps you to balance the auditing activity.

## Deciding Who and What to Audit

Novice system administrators tend to start out wanting to collect as much information as possible about user and system actions. However, audit files can quickly grow to fill the available space, so it makes sense to be selective about what kinds of activities are audited.

Issue	Description
1. Determine which audit classes you need for your site	The best time to add audit classes or to change the default classes is before you start the auditing service. See "Audit Flags" on page 368 for information about auditing classes.
2. Determine event-to-class mappings	In many situations, the default mapping is sufficient. However, if you have added new classes or have changed them, you might also need to move an event to a different class.
3. Decide what classes should be audited for all users on all machines	The system-wide audit flags in the <code>audit_control</code> file apply to all users and processes. Audit flags determine whether an audit class is audited for success, for failure, and for both. Every machine in the installation should have the same audit flag settings in their <code>audit_control</code> file.
4. Determine user exceptions to the installation-wide audit settings	If you decide that some users should be audited differently from the system-wide settings, modify the users' entries in the <code>audit_user</code> file on each machine. For more information see "How to Change Users' Audit Characteristics" on page 340.

Issue	Description
5. Determine the minimum free disk space ( <code>minfree</code> ) that should be available on an audit file system before a warning is sent	When the amount of disk space on an audit file system drops below the <code>minfree</code> percentage, the audit daemon switches to the next available audit directory and sends a warning that the soft limit has been exceeded. For more information, see “Example — Changing the Soft Limit for Warnings” on page 339.
6. Decide which audit policies your site needs	The policy variable is a dynamic kernel variable, so its value is not saved when the system is brought down. Therefore, you should set the desired policy by using an appropriate startup script. For more information, see “How to Enable or Disable an Audit Policy” on page 346.
7. Decide how to manage the <code>audit_warn</code> mail alias	The <code>audit_warn</code> script is run by the <code>auditd</code> daemon whenever the daemon switches audit directories or encounters difficulty (such as lack of disk space). By default, the <code>audit_warn</code> script sends mail to an <code>audit_warn</code> alias and sends a message to the console. You need to either modify the alias, or change the script to send the message to a different alias. For more information, see “How to Configure the <code>audit_warn</code> Alias” on page 346.
8. Decide what to do when all the audit directories are full	To permit the system to continue functioning in the event of an audit trail overflow, you can enable the <code>cnt</code> policy. For a <code>cnt</code> policy example, see “Example — Setting the <code>cnt</code> Policy” on page 347.  Alternatively, you can create an account that can log in and work without being audited. See “Example — Creating an Audit Admin Login” on page 341 for an audit account example.

## Determining Which Audit Policies to Use

Audit policies refer to the auditing options that are enabled or disabled for a particular configuration. You can inspect, enable, or disable the current audit policy with the `auditon` system call at the program level. Or, to do the same task you can run the `auditconfig` command.

By default, all audit policies are disabled. You need to enable any audit policies that you want to use.

The audit policies are disabled by default to minimize storage requirements and system processing demands. You can enable audit policies, and disable them dynamically. Use the following table to determine if the needs of your site justify the additional overhead that results from enabling one or more audit policies.

Policy Name	Description	Why Change the Policy?
arge	<p>When disabled, this policy omits environment variables of an executed program script from the <code>exec</code> audit record.</p> <p>When enabled, this policy adds the environment variables of an executed program script to the <code>exec</code> audit record. The resulting audit records contain much more detail than when this policy is disabled.</p>	<p>The disabled option collects much less information than the enabled option.</p> <p>The enabled option makes sense when you are auditing a few users, or when you have suspicions about the environment variables that are being used in <code>exec</code> programs.</p>
argv	<p>When disabled, this policy omits the arguments of an executed program script from the <code>exec</code> audit record.</p> <p>When enabled, this policy adds the arguments of an executed program script to the <code>exec</code> audit record. The resulting audit records contain much more detail than when this policy is disabled.</p>	<p>The disabled option collects much less information than the enabled option.</p> <p>The enabled option makes sense when you are auditing a few users, or when you have reason to believe that unusual <code>exec</code> programs are being run.</p>
cnt	<p>When disabled, this policy blocks a user or application when audit records can not be added to the audit trail because no disk space is available.</p> <p>When enabled, this policy allows the event to complete without an audit record being generated. A count of audit records that are dropped is maintained.</p>	<p>The disabled option makes sense in an environment where security is paramount.</p> <p>The enabled option makes sense when system availability is more important than security.</p>
group	<p>When disabled, this policy does not add a groups list to audit records.</p> <p>When enabled, this policy adds a groups list to every audit record as a special token.</p>	<p>The disabled option usually satisfies requirements for site security.</p> <p>The enabled option makes sense when you need to audit which groups are generating auditable events.</p>
path	<p>When disabled, this policy records in an audit record at most one path that is used during a system call.</p> <p>When enabled, this policy records every path that is used in conjunction with an audit event to every audit record.</p>	<p>The disabled option places at most one path in an audit record.</p> <p>The enabled option enters each file name or path that is used during a system call in the audit record as a path token.</p>
seq	<p>When disabled, this policy does not number the audit records in sequence.</p> <p>When enabled, this policy adds a sequence number (<code>seq</code> token) to every audit record.</p>	<p>The disabled option is sufficient when auditing is running smoothly.</p> <p>The enabled option makes sense when you are checking that audit files are being written correctly. In the case of file corruption (for example, a partially written audit record), you may be able to spot bad records faster if the sequence numbers are out of order or if some numbers are missing.</p>

Policy Name	Description	Why Change the Policy?
trail	When disabled, this policy does not add a trailer token to audit records.  When enabled, this policy adds a trailer token to every audit record.	The disabled option creates a smaller audit record.  The enabled option marks the end of each audit record clearly with a trailer token. The trailer token is often used in conjunction with the sequence token when debugging. In the case of file corruption (for example, a partially written audit record), the <code>auditreduce</code> command resyncs faster on good records.

## Controlling Auditing Costs

Because auditing consumes system resources, you must control the degree of detail that is recorded. When you decide what to audit, consider the following costs of auditing:

- Cost of increased processing time
- Cost of analysis of audit data
- Cost of storage of audit data

### Cost of Increased Processing Time

The cost of increased processing time is the least significant of the costs of auditing. The first reason is that auditing generally does not occur during computation-intensive tasks, such as image processing, complex calculations, and so forth. The other reason that processing cost is usually insignificant is that the cost for single-user systems is usually small enough to ignore.

### Cost of Analysis

The cost of analysis is roughly proportional to the amount of audit data that is collected. The cost of analysis includes the time it takes to merge and review audit records, and the time it takes to archive them and keep them in a safe place.

The fewer records that you generate, the less time it takes to analyze them. Upcoming sections, “Cost of Storage” on page 334 and “Auditing Efficiently” on page 334, describe how you can reduce the amount of data that you collect, while still providing enough coverage to achieve your site’s security goals.

## Cost of Storage

Storage cost is the most significant cost of auditing. The amount of audit data depends on the following:

- Number of users
- Number of machines
- Amount of use
- Degree of security that is required

Because these factors vary from site to site, no formula can determine in advance the amount of disk space to set aside for audit data storage.

Full auditing (with the `all` flag) fills up disks quickly. Even a simple task such as compiling a program of modest size (for example, 5 files, 5000 lines total) in less than a minute could generate thousands of audit records, occupying many megabytes of disk space. Therefore, it is very important to use the preselection features to reduce the volume of records that are generated. For example, by omitting the `fr` class instead of all classes, you can reduce the audit volume by more than two-thirds. Efficient audit file management is also important after the audit records are created, to reduce the amount of storage that is required.

Before you configure auditing, you should understand the audit flags and the types of events they flag. Develop a philosophy of auditing for your site that is based on the amount of security your site requires, and the types of users you administer.

---

## Auditing Efficiently

The techniques in this section can help you achieve your organization's security goals while auditing more efficiently:

- Randomly audit only a certain percentage of users at any one time
- Reduce the disk-storage requirements for audit files by combining, reducing, and compressing them. Develop procedures for archiving them, transferring them to removable media, and storing them offline.
- Monitor the audit data in real time for unusual behaviors. You can set up procedures to monitor the audit trail as it is generated for certain activities. You can write a script to trigger an automatic increase in the auditing of certain users or certain machines in response to detection of unusual events.

For example, write a script that (1) monitors the creation of audit files on all the audit file servers, and (2) processes them with the `tail` command (see the `tail(1)` man page). Pipe the output of `tail -of` through the `praudit` command to yield a stream of audit records as soon as the records are generated. This stream can be analyzed for unusual message types or other indicators, and delivered to the

auditor. Or, the script can be used to trigger automatic responses.

In addition, the script should include code that (3) constantly monitors the audit directories for the appearance of new `not_terminated` audit files, and (4) terminates outstanding `tail` processes when their files are no longer being written to.





## Managing the BSM Service (Tasks)

---

This chapter presents procedures that are designed to help you set up and manage a Solaris environment that includes auditing. This chapter also includes instructions for administering the audit trail and for administering device allocation. This is a list of the task maps in this chapter.

- “Managing BSM (Task Map)” on page 337
- “Configuring Audit Files (Task Map)” on page 338
- “Configuring the Audit Service (Task Map)” on page 343
- “Managing Audit Records (Task Map)” on page 349
- “Adding an Allocatable Device (Task Map)” on page 352

---

### Managing BSM (Task Map)

The following task map shows the major tasks that are required to administer the BSM services.

Task	Description	For Instructions
Plan for auditing	Configuration issues to consider and make decisions about, before you configure auditing.	Chapter 23
Configure audit files	Defines which events, classes, and users require auditing.	“Configuring Audit Files” on page 338
Configure auditing	Configures each host so you can use auditing.	“Configuring the Audit Service” on page 343
Manage audit records	Merges and analyzes the audit data.	“Managing Audit Records” on page 348

Task	Description	For Instructions
Manage device allocation	Defines which devices should be accessed through the device allocation mechanism.	"Managing Device Allocation" on page 352

## Configuring Audit Files

Before you enable auditing on your network, you may want to edit the audit configuration files. Many of the following procedures require that you restart the service or reboot the local system. You should make as many of these changes as possible before you start the service.

### Configuring Audit Files (Task Map)

The following task map describes the tasks in this section.

Task	Description	For Instructions
Change audit flags	Defines the location of the audit directories and system-wide flags for the audit service.	"How to Change Audit Flags" on page 338
Change audit characteristics for users	Selects specific auditing for a user.	"How to Change Users' Audit Characteristics" on page 340
Change audit classes	Selects which events, classes, and users require auditing.	"How to Change Audit Classes" on page 341
Change audit events	Adds new events to the auditing service.	"How to Change Audit Events" on page 342

#### ▼ How to Change Audit Flags

Audit flags are defined in the `/etc/security/audit_control` file. The audit flags select which classes of audit records are written to the audit log.

1. **Become superuser or assume an equivalent role.**
2. **(Optional) Save a backup copy of the `audit_control` file.**

```
# cp /etc/security/audit_control /etc/security/audit_control.save
```

3. **Add new entries to the `audit_control` file.**

Each entry has the following format:

*title*:*string*

*title* Defines the type of line. Options are `dir:`, `flags:`, `minfree:`, or `naflags:`.

*string* Lists specific data that is associated with the line type

#### 4. Instruct the audit daemon to read the new `audit_control` file.

The audit daemon stores the information internally. To use the new information, either reboot the system or type the following command:

```
# audit -s
```

### Example — Changing Audit Trail File Locations

Lines that start with `dir:` define which audit file systems can be used to store audit trail files. In this example, two additional locations for audit trail files are defined.

```
# cat /etc/security/audit_control
dir:/etc/security/audit/host.1/files
dir:/etc/security/audit/host.2/files
dir:/var/audit
flags:
minfree:10
naflags:lo
```

### Example — Changing Audit Flags for All Users

The `flags` line in the `audit_control` file defines which classes of events are audited for all users on the host. The classes are separated by commas, with no spaces. In this example, the events in the `lo` class are audited for all users.

```
# cat /etc/security/audit_control
dir:/var/audit
flags:lo
minfree:10
naflags:lo
```

### Example — Changing the Soft Limit for Warnings

The `minfree` line in the `audit_control` file defines the minimum free-space level for all audit file systems. In this example, the soft limit is set so that a warning is issued when only 10 percent of the file system is available.

```
# cat /etc/security/audit_control
dir:/var/audit
flags:
```

```
minfree:10
naflags:lo
```

## Example — Changing Auditing of Nonattributable Events

The `naflags:` line in the `audit_control` file defines which classes of nonattributable events are audited for all users on the host. The classes are separated by commas, with no spaces. In this example, the `na` event class was added.

```
# cat /etc/security/audit_control
dir:/var/audit
flags:
minfree:10
naflags:lo,na
```

## ▼ How to Change Users' Audit Characteristics

Definitions for each user can be stored in the `/etc/security/audit_user` file.

1. **Become superuser or assume an equivalent role.**
2. **(Optional) Save a backup copy of the `audit_user` file.**

```
# cp /etc/security/audit_user /etc/security/audit_user.save
```

3. **Add new entries to the `audit_user` file.**

Each entry has the following format:

```
username:always:never
```

<i>username</i>	Selects the name of the user to be audited
<i>always</i>	Selects the list of audit classes that should always be audited
<i>never</i>	Selects the list of audit classes that should never be audited

You can specify multiple flags by separating the audit classes with commas. For more information about audit flags, see “Audit Flags” on page 368.

4. **Make the new data available to the BSM service.**

To use the new data, either reboot the system, or have the user log out and back in again.

## Example — Changing Auditing for One User

This example shows an entry that causes audit records to be generated anytime the user `sue` accesses any programs in the `login` class (`lo`).

```
# grep sue /etc/security/audit_user
sue:lo:
```

## Example — Creating an Audit Admin Login

If all the audit partitions are full, then it could be impossible to log in to a host. If all logins are audited, then the fact that the audit partitions are full would prevent anyone from completing a login. To avoid this situation, you can set up a special login that is not audited. This new login would allow you to log in to the host even if the audit partitions are full. Then, you could fix the problem with the full partitions. In this example, the user `auditadm` is defined so that no auditing takes place.

```
# grep auditadm /etc/security/audit_user
auditadmin:no:yes
```

---

**Note** – The user login that is selected to serve as the audit admin login might need to be monitored in another way.

---

## ▼ How to Change Audit Classes

Audit classes are defined in the `/etc/security/audit_class` file.

1. **Become superuser or assume an equivalent role.**
2. **(Optional) Save a backup copy of the `audit_class` file.**

```
# cp /etc/security/audit_class /etc/security/audit_class.save
```

3. **Add new entries to the `audit_class` file.**

Each entry has the following format:

```
0xnumber : name : description
```

<i>number</i>	Defines the unique audit class mask
<i>name</i>	Defines the two-letter name of the audit class
<i>description</i>	Defines the descriptive name of the audit class

4. **Make the new data available to the BSM service.**

To use the new data, either reboot the system, or type the following command:

```
# auditconfig -conf
```

## Example — Setting a New Audit Class

In step 3, add an entry that resembles the following to set a new audit class called `de`:

```
0x00010000:de:device allocation
```

## ▼ How to Change Audit Events

Audit event definitions are stored in the `/etc/security/audit_event` file. A record is generated only after the event definition has been created and a user-level action generates the event.

1. **Become superuser or assume an equivalent role.**
2. **(Optional) Save a backup copy of the `audit_event` file.**

```
# cp /etc/security/audit_event /etc/security/audit_event.save
```

3. **Add new entries to the `audit_event` file.**

Each entry has the following format:

```
number : name : description : classes
```

<i>number</i>	Defines a unique audit event number, which must start after 32768.
<i>name</i>	Defines the unique audit event name.
<i>description</i>	Describes the audit event. Often includes the name of the man page for the audit event
<i>classes</i>	Selects the audit classes that include this event.

4. **Make the new data available to the BSM service.**

To use the new data, either reboot the system, or type the following command:

```
# auditconfig -conf
```

## Example — Adding a New Audit Event

This example shows an entry that defines a new audit event for a local application.

```
# grep localapp /etc/security/audit_event
32769:aue_localapp:localapp(1):ap
```

---

# Configuring the Audit Service

This section covers the tasks that are required to configure and enable the audit service.

## Configuring the Audit Service (Task Map)

The following task map describes the tasks that are required to configure auditing.

Task	Description	For Instructions
1. Plan for auditing	Resolve configuration issues before you configure auditing.	Chapter 23
2. Create audit partitions	Creates the partitions for the audit files.	"How to Create Partitions for Auditing" on page 343
3. Create the <code>audit_warn</code> alias	Defines who should get email warnings.	"How to Configure the <code>audit_warn</code> Alias" on page 346
4. (Optional) Change audit policies	Defines additional audit records or auditing conditions.	"How to Enable or Disable an Audit Policy" on page 346
5. (Optional) Change the audit configuration files	Selects which events, classes, and users require auditing.	"Configuring Audit Files" on page 338
6. Enable auditing	Turns on auditing.	"How to Enable Auditing" on page 347
7. (Optional) Disable auditing	Turns off auditing.	"How to Disable Auditing" on page 348
8. (Optional) Start device allocation	Selects which removable media should be accessed in a more secure mode.	"Managing Device Allocation" on page 352

### ▼ How to Create Partitions for Auditing

The following procedure shows how to create partitions for auditing, as well as the corresponding file systems and directories. Skip steps as necessary, depending on if you already have an empty partition, or if you have already mounted an empty file system.

1. **Become superuser or assume an equivalent role.**
2. **Determine the amount of disk space that is required.**

Assign at least 200 Mbytes of disk space per host. However, the disk space requirements are based on how much auditing you perform. So, your requirements might be far greater than this figure. Remember to include a partition for a directory of last resort.

**3. Create dedicated audit partitions, as needed.**

This step is most easily done during server installation. You can also create the partitions on disks that have not yet been mounted on the server. For complete instructions on how to create the partitions, see “Creating a UFS File System” in *System Administration Guide: Basic Administration*.

```
newfs /dev/rdisk/cwtxdysz
```

Where `/dev/rdisk/cwtxdysz` is the raw device name for the partition.

If the local host is to be audited, create an audit directory of last resort for it as well.

**4. Create mount points for each new partition.**

```
mkdir /var/audit/server-name.n
```

Where `server-name.n` is the name of the server and a number that identifies each partition. The number is optional, but the number is useful when there are many audit directories.

**5. Add entries to automatically mount the new partitions.**

Add a line to the `/etc/vfstab` file that resembles the following:

```
/dev/dsk/cwtxdysz /dev/rdisk/cwtxdysz /var/audit/server-name.n ufs 2 yes
```

**6. (Optional) Remove the minimum free space threshold on each partition.**

If you use the default configuration, a warning will be generated when the directory is 80 percent full, so there is no reason to reserve free space on the partition.

```
tunefs -m 0 /var/audit/server-name.n
```

**7. Mount the new audit partitions.**

```
mount /var/audit/server-name.n
```

**8. Create audit directories on the new partitions.**

```
mkdir /var/audit/server-name.n/files
```

**9. Correct the permissions on the mount points and new directories.**

```
chmod -R 750 /var/audit/server-name.n/files
```

**10. (Optional) On a file server, define the file systems to be made available to other hosts.**

Often, disk farms are installed to store the audit records. If an audit directory is to be used by several systems, then the directory must be shared through the NFS service. Add an entry resembling the following for each directory to the `/etc/dfs/dfstab` file.



```
share -F nfs /var/audit/server-name.n/files
```

### 11. (Optional) On a file server, restart the NFS service.

If this command the first `share` command or set of `share` commands that you have initiated, it is probable that the NFS daemons are not running. The following commands kill the daemons and restart them. Refer to “Setting Up NFS Services” in *System Administration Guide: Resource Management and Network Services* for more information about the NFS service.

```
# /etc/init.d/nfs.server stop
# /etc/init.d/nfs.server start
```

## Example — Creating an Audit Directory of Last Resort

All systems that run the auditing service should have a local file system that can be used if no other file system is available. In this example, a file system is being added to a system named `egret`. Since this file system is only used locally, none of the steps for a file server are followed.

```
# newfs /dev/rdisk/c0t2d0
# mkdir /var/audit/egret
# grep egret /etc/vfstab
/dev/dsk/c0t2d0s1 /dev/rdisk/c0t2d0s1 /var/audit/egret ufs 2 yes -
# tunefs -m 0 /var/audit/egret
# mount /var/audit/egret
# mkdir /var/audit/egret/files
# chmod -R 750 /var/audit/egret/files
```

## Example — Creating New Audit Partitions

In this example, a new file system is created on two new disks that are to be used by other systems in the network.

```
# newfs /dev/rdisk/c0t2d0
# newfs /dev/rdisk/c0t2d1
# mkdir /var/audit/egret.1
# mkdir /var/audit/egret.2
# grep egret /etc/vfstab
/dev/dsk/c0t2d0s1 /dev/rdisk/c0t2d0s1 /var/audit/egret.1 ufs 2 yes -
/dev/dsk/c0t2d1s1 /dev/rdisk/c0t2d1s1 /var/audit/egret.2 ufs 2 yes -
# tunefs -m 0 /var/audit/egret.1
# tunefs -m 0 /var/audit/egret.2
# mount /var/audit/egret.1
# mount /var/audit/egret.2
# mkdir /var/audit/egret.1/files
# mkdir /var/audit/egret.2/files
# chmod -R 750 /var/audit/egret.1/files /var/audit/egret.2/files
# grep egret /etc/dfs/dfstab
share -F nfs /var/audit/egret.1/files
share -F nfs /var/audit/egret.2/files
```

```
# /etc/init.d/nfs.server stop
# /etc/init.d/nfs.server start
```

## ▼ How to Configure the `audit_warn` Alias

The `audit_warn` script generates mail to an alias called `audit_warn`. To send this mail to a valid email address, you can follow either of the following steps:

1. **Become superuser or assume an equivalent role.**
2. **(Optional) Swap the `audit_warn` alias with another alias.**

One option is to edit the `audit_warn` script and replace `audit_warn` with another alias. After you swap `audit_warn` for `root`, the line that sends the email message would resemble the following:

```
/usr/ucb/mail -s "$SUBJECT" root
```

Ten lines in the script require this change.

3. **(Optional) Redirect the `audit_warn` email to another alias.**

The other option is to redirect the email in the `/etc/mail/aliases` file. In this case, you would add an alias similar to the following to the local `/etc/mail/aliases` file or to the `mail_aliases` database in the name space. The new entry would resemble the following if the email were to be redirected to the `root` alias:

```
audit_warn: root
```

## ▼ How to Enable or Disable an Audit Policy

Audit policies determine the characteristics of the audit records for the local host. Audit policies are either enabled or disabled for a particular configuration. By default, all audit policies are disabled. You need to enable any audit policies that you want to use. For a description of each policy, see “Audit Policies” on page 371.

1. **Become superuser or assume an equivalent role.**
2. **(Optional) Review the existing audit policies.**

Ensure that you are aware of all the policies that are being used before you change any. The following command lists the enabled policies:

```
# auditconfig -lspolicy
```

3. **Enable or disable the audit policy.**

```
auditconfig -setpolicy flagpolicyname
```

<i>flag</i>	A + enables the policy. A - disables the policy
<i>policyname</i>	Selects the policy to be enabled or disabled

The policy is in effect until the next boot, or until the policy is modified by the `auditconfig-setpolicy` command.

## Example — Setting the `cnt` Policy

The `cnt` policy can be set so that if the audit partitions become full, then processes are not blocked. The records are discarded when the partitions are full, but the system still functions even though the auditing process is not recording the events. The `cnt` policy should not be set if security is paramount, since unrecorded events can occur if the file system is full.

The following command enables the `cnt` policy:

```
# auditconfig -setpolicy +cnt
```

For a secure site, you should enable the `cnt` policy in an appropriate startup file.

## ▼ How to Enable Auditing

This task starts the auditing service. If the service has been configured, then rebooting the host also starts the service.

1. **Become superuser or assume an equivalent role.**
2. **Bring the system into single-user mode.**

```
# /etc/telinit 1
```

See the `telinit(1M)` man page for more information.

3. **Run the script to configure the system to run auditing.**

Go to the `/etc/security` directory, and execute the `bsmconv` script there. The script sets up a standard Solaris machine to run BSM after a reboot. See the `bsmconv(1M)` man page.

```
# cd /etc/security
# ./bsmconv
```

4. **Bring the system into multiuser mode.**

```
# /etc/telinit 6
```

The startup file `/etc/security/audit_startup` causes the audit daemon to run automatically when the system enters multiuser mode.

---

**Note** – The `bsmconv` script adds a line to the `/etc/system` file that prevents users from aborting the system with the Stop-A keyboard sequence. To retain the ability to abort the system with the Stop-A keyboard sequence, you must comment out the line in the `/etc/system` file that reads: `set abort_enable=0`.

---

## ▼ How to Disable Auditing

If BSM is no longer required at some point, you can disable it by running the `bsmunconv` command. See the `bsmconv(1M)` man page.

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

**2. Bring the system into single-user mode.**

```
# /etc/telinit 1
```

See the `telinit(1M)` man page for more information.

**3. Run the script to disable auditing.**

Change to the `/etc/security` directory, and execute the `bsmunconv` script there.

```
# cd /etc/security
# ./bsmunconv
```

**4. Bring the system into multiuser mode.**

```
# /etc/telinit 6
```

---

**Note** – The `bsmunconv` script removes the line in the `/etc/system` file that allows users to abort the system with the Stop-A keyboard sequence. If you want to continue to prevent users from aborting the system with the Stop-A keyboard sequence after you run the `bsmunconv` script, you must reenter into the `/etc/system` file the line that reads: `set abort_enable=0`.

---

---

## Managing Audit Records

Managing the audit trail enables you to monitor the actions of users on your network. Auditing can generate large amounts of data. The following tasks show you how to work with all this data.

## Managing Audit Records (Task Map)

The following task map describes the tasks in this section.

Task	Description	For Instructions
Merge audit records	Combines audit files from several machines into one audit trail.	“How to Merge Audit Records” on page 349
Display audit record formats	Displays the order of tokens for a particular audit event.	“How to Display Audit Record Formats” on page 351
Prevent audit trail overflow	Prevents the audit file systems from completely filling up.	“How to Prevent Audit Trail Overflow” on page 351

### ▼ How to Merge Audit Records

This task shows you how to merge all audit files in all the audit directories. Follow these steps when you want to analyze the contents of the audit trail.

1. **Become superuser or assume an equivalent role.**
2. **Change directories to the primary audit directory.**

```
# cd /etc/security/audit/server-name.1/files
```

Moving to this directory places the merged file in this protected directory.

3. **Merge the audit records.**

```
# auditreduce > merged.log
```

All directories that are listed in the `dir:` lines of the `audit_control` file on *server-name* are merged and placed in the file called `merged.log`.

### Example — Displaying the Entire Audit Trail

To display the entire audit trail at once, pipe the output of the `auditreduce` command into the `praudit` command.

```
# auditreduce | praudit
```

### Example — Printing the Entire Audit Trail

With a pipe to the `lp` command, the output goes to the printer.

```
# auditreduce | praudit | lp
```

## Example — Combining and Reducing Audit Files

Use `auditreduce` with the `-O` option to combine several audit files into one file and to save them in a specified output file. `auditreduce` can do this type of combination and deletion automatically (see the `-C` and `-D` options in the `auditreduce(1M)` man page). However, it is often easier to select the files manually (perhaps with the `find` command) and use `auditreduce` to combine just the named set of files.

When used in this way, `auditreduce` merges all the records from its input files into a single output file. The input files should then be deleted. In addition, the output file should be kept in a directory that is named `/etc/security/audit/server-name/files` so that `auditreduce` can find it.

```
# auditreduce -O combined-filename
```

The `auditreduce` command can also reduce the number of records in its output file by eliminating the less interesting records as it combines the input files. For example, you might use `auditreduce` to retain only the login and logout records in audit files that are over a month old. If you need to retrieve the complete audit trail, you could recover it from backup tapes.

```
# auditreduce -O daily.summary -b 19990413 -c lo; compress *daily.summary
# mv *daily.summary /etc/security/summary.dir
```

## Example — Displaying User Activity From a Selected Date

In the following example, the system administrator checks to see when user `tamiko` logged in and logged out on April 13, 1999, by requesting the `lo` event class. The short-form date is in the form `yyymmdd`. The long form is described in the `auditreduce(1M)` man page.

```
# auditreduce -d 990413 -u tamiko -c lo | praudit
```

## Example — Copying Selected Records to a Single File

In this example, login and logout messages for a particular day are selected from the audit trail and merged into a target file. The target file is written in a directory other than the normal audit root directory.

```
# auditreduce -c lo -d 990413 -O /usr/audit_summary/logins
```

The `-O` option creates an audit file with 14-character timestamps for both the start-time and the end-time, with the suffix `logins`:

```
/usr/audit_summary/19990413000000.19990413235959.logins
```

## Example — Cleaning Up a not\_terminated Audit File

Occasionally, an audit daemon dies while its audit file is still open, or a server becomes inaccessible and forces the machine to switch to a new server. In such instances, an audit file remains with the string `not_terminated` as the end-time, even though the file is no longer used for audit records. When you find such a file, you can manually verify that the file is no longer in use and clean it up by specifying the name of the file with the correct options.

```
# audit -s
19990414121112.not_terminated.egret
# auditreduce -O egret 19990413120429.not_terminated.egret
```

The `audit` command checks the name of the current audit file. The `auditreduce` command creates a new audit file with the correct name and correct timestamps, with the correct suffix (`egret`), and copies all the records into it.

## ▼ How to Display Audit Record Formats

The following command displays the formats of all audit event records. This command operates on records in the `audit_class`, `audit_event`, and `audit_record_attr` files.

```
# bsmrecord -a
```

## Example — Displaying the Format of an Audit Record

In this example, the format of the audit record with an ID of 6152 is displayed.

```
# bsmrecord -i 6152
# login: terminal login
program /usr/sbin/login see login(1)
event ID 6152 AUE_login
class lo (0x00001000)
header-token
subject-token
text-token error message
exit-token
#
```

## ▼ How to Prevent Audit Trail Overflow

If your security policy requires that all audit data be saved, do the following:

1. **Set up a schedule to regularly archive audit files and to delete the archived audit files from the audit file system.**

2. **Manually archive audit files by backing them up on tape or by moving them to an archive file system.**
3. **Store context-sensitive information that will be needed to interpret audit records, along with the audit trail.**
4. **Keep records of which audit files are moved offline.**
5. **Store the archived tapes appropriately.**
6. **Reduce the volume of audit data that you store by creating summary files.**

You can extract summary files from the audit trail by using options to `auditreduce` so that the summary files contain only records for certain specified types of audit events. For examples, see “Example — Combining and Reducing Audit Files” on page 350 and “Example — Copying Selected Records to a Single File” on page 350.

---

## Managing Device Allocation

You can use device allocation to decrease the security risk that is associated with various removable media.

### Adding an Allocatable Device (Task Map)

The following task map describes the major steps that are required to define a new allocatable device.

Task	Description	For Instructions
1. Create or change an entry in the <code>device_allocate</code> file	Defines which devices are controlled by the device-allocation mechanism.	“How to Change Which Devices Can Be Allocated” on page 353
2. Create a lock file	Enables the device allocation mechanism to work on a specific device.	“How to Set Up Lock Files for an Allocatable Device” on page 353
3. (Optional) Create a device-clean script	Purges data from a physical device.	“Device-Clean Scripts” on page 396
4. Allocate the device	Adds a device to the device-allocation mechanism.	“How to Allocate a Device” on page 354
5. (Optional) Deallocate the device	Removes a device from use.	“How to Deallocate a Device” on page 354



## ▼ How to Set Up Lock Files for an Allocatable Device

The lock files are zero-length files that are created in the `/etc/security/dev` directory. One file is created for each allocatable device. If no lock file exists for a device, the device cannot be allocated, so no one can access the device.

1. **Become superuser or assume an equivalent role.**
2. **Obtain the device name for the device from its entry in the `device_maps` file by using the `dminfo` command.**

See “The `device_maps` File” on page 393 and the `dminfo(1M)` and `device_maps(4)` man pages. For example, the device name for device type `st` is `st0`. In the next step, you will use the device name as the name of the lock file.
3. **Create an empty lock file for the device by using the `touch` command.**

Use the device name for the file name in place of *device-name*.

```
# cd /etc/security/dev
# touch device-name
# chmod 600 device-name
# chown bin device-name
# chgrp bin device-name
```

## ▼ How to Change Which Devices Can Be Allocated

This procedure defines which devices can be used with the device allocation mechanism.

1. **Become superuser or assume an equivalent role.**
2. **Determine which devices are listed in the `/etc/security/device_allocate` file.**
3. **Decide if there are devices that are not in the `device_allocate` file, yet should be made allocatable.**
4. **Edit the `device_allocate` file and add the new device.**

Each entry should use the following format:

```
device-name ; device-type ; ; ; program
```

<i>device-name</i>	Specifies the name of the device
<i>device-type</i>	Specifies the device type
<i>program</i>	Specifies the purge program to be run

## ▼ How to Allocate a Device

1. Become superuser or assume an equivalent role.
2. Use the `allocate` command with a device that is specified by device name.

```
sar1% allocate st0
```

You can also allocate a device by device type by using the `-g` option to the `allocate` command.

If the command cannot allocate the device, an error message is displayed in the console window. For a list of allocation error messages, see the `allocate(1)` man page.

### Example — Allocating a Printer

```
sar1% allocate /dev/lp/chestnut
```

Only the user who ran the `allocate` command can use the printer.

## ▼ How to Deallocate a Device

Deallocate a device by using the `deallocate` command followed by the device file name.

```
sar1% deallocate st0
```

Deallocation enables other users to allocate and use the device when you are finished.

### Example — Deallocating a Printer

To deallocate a printer that is named `chestnut`, type the following command:

```
# deallocate /dev/lp/chestnut
```

### Example — Forcing a Deallocation

Devices that a user has allocated are not automatically deallocated when the process terminates or when this user logs out. You most commonly need to use the following form of the `deallocate` command when a user forgets to deallocate a specific device. Then you force its deallocation, so that others users can allocate it.

```
# deallocate -F st0
```

## Example — Deallocating All Devices



---

**Caution** – You can deallocate all devices only at system initialization time.

---

```
# deallocate -I
```



## BSM Services Reference

---

This chapter describes the important components of the BSM module, auditing and device allocation.

Auditing helps you detect potential security breaches by revealing suspicious or abnormal patterns of system usage. Auditing also provides a means to trace suspect actions back to a particular user, thus serving as a deterrent. If users know that their activities are likely to be audited, they might be less likely to attempt malicious activities.

This is a list of the reference information in this chapter.

- “Audit Commands” on page 357
- “Audit Files” on page 363
- “Audit Flags” on page 368
- “Audit Policies” on page 371
- “Process Audit Characteristics” on page 371
- “Audit Trail” on page 372
- “Audit Record Structure” on page 374
- “Audit Token Formats” on page 374
- “Device Allocation Reference” on page 392

---

### Audit Commands

This section provides information about the commands that are used with the auditing service.

### The Audit Daemon

The following list summarizes what the audit daemon, `auditd`, does.

- `auditd` opens and closes audit log files in the directories that are specified in the `audit_control` file, in the order in which they are specified.
- `auditd` reads audit data from the kernel and writes it to an audit log file.
- `auditd` executes the `audit_warn` script when the audit directories fill past limits that are specified in the `audit_control` file. The script, by default, sends warnings to the `audit_warn` mail alias and to the console.
- By default, when all audit directories are full, processes that generate audit records are suspended. In addition, the `auditd` command writes a message to the console and to the `audit_warn` mail alias. The audit policy can be reconfigured with the `auditconfig`. At this point, only the system administrator can log in to write audit files to tape, delete audit files from the system, or do other cleanup.

The `auditd` daemon can be started automatically when the machine is brought up to multiuser mode, or you can start it from the command line. When the audit daemon is started, it determines the amount of free space necessary for audit log files.

The daemon uses the list of audit directories in the `audit_control` file as possible locations for creating audit files. The audit daemon maintains a pointer into this list of directories, starting with the first directory. Every time the audit daemon needs to create an audit file, it puts the file into the first available directory in the list, starting at the audit daemon's current pointer. You can reset the pointer to the beginning of the list by running the `audit -s` command. When you use the `audit -n` command to instruct the daemon to switch to a new audit file, the new file is created in the same directory as the current file.

## The `audit` Command

The `audit` command controls the actions of the audit daemon. The `audit` command can do the following tasks:

- Enable and disable auditing
- Reset the audit daemon
- Adjust the auditing preselection mask on the local machine
- Write audit records to a different audit log file

See the `audit(1M)` man page for a discussion of the available options.

## The `auditreduce` Command

Use the `auditreduce` command to merge audit records from one or more input audit files or to perform a post selection of audit records. See the `auditreduce(1M)` man page. To merge the entire audit trail, you can run this command on the machine on which all the audit file systems for the installation are mounted.

The `auditreduce` command enables you to track all auditable actions on multiple machines from a single location. If you identically configure all machines at an installation for auditing, and create servers and local directories for the audit log files, then the `auditreduce` command can read the logical combination of all audit files in the installation as a single audit trail. `auditreduce` ignores how the records were generated or where they are stored. Without options, the `auditreduce` command merges audit records from all the audit files in all of the subdirectories in the audit root directory (`/etc/security/audit`) and sends the result to standard output. You can also place it into a single, chronologically ordered output file. The file contains binary data.


The `auditreduce` command also can select particular types of records for analysis. The merging and selecting functions of `auditreduce` are logically independent. `auditreduce` captures data from the input files as the records are read, before the files are merged and written to disk.

The `praudit` command makes the binary output of `auditreduce` readable.


By specifying options to the `auditreduce` command, you can also do the following:

- Request audit records that were generated by only certain audit flags
- Request audit records that were generated by one particular user
- Request audit records that were generated on specific dates

With no arguments, `auditreduce` checks the subdirectories within the `/etc/security/audit` directory, the default audit root directory. The command checks for a `files` directory in which the `start-time.end-time.hostname` files reside. The `auditreduce` command is very useful when the audit data for different hosts (Figure 25–1) or for different audit servers (Figure 25–2) reside in separate directories.



**FIGURE 25–1** Audit Trail Storage Sorted by Host



**FIGURE 25-2** Audit Trail Storage Sorted by Server

When you do not store audit data in the default directory, perhaps because the partition for `/etc/security/audit` is very small, you can pass the `auditreduce` command another directory by using the `-R` option:

```
# auditreduce -R /var/audit-alt
```

You can also specify a particular subdirectory by using the `-S` option:

```
# auditreduce -S /var/audit-alt/host1
```

You can direct `auditreduce` to process only certain audit log files by specifying them as command arguments:

```
# auditreduce /var/audit/egret/files/2001*.2001*egret
```

For other options and additional examples, see the `auditreduce(1M)` man page.

## The `praudit` Command

The `praudit` command reads audit records in binary format from standard input and displays the records in a presentable format. The input can be piped from the `auditreduce` command or from a single audit file. Input can also be produced with the `cat` command to concatenate several files, or the `tail` command for a current audit file.

The `praudit` command can generate four output formats:

- Default** – The default option displays one audit token per line. The default option displays the audit event by its description, such as `ioctl(2)`, and displays any value that could be text in text format. For example, a user is displayed as the user name, not as the user ID.



- **-l option** – The long option displays one audit record per line. The `-d` option changes the delimiter used between token fields, and between tokens. The default delimiter is a comma.
- **-r option** – The raw option displays any value that could be numeric as a number. For example, a user is displayed by user ID, Internet addresses are in hexadecimal format, and modes are in octal format. The audit event is displayed as its event number, such as 158.
- **-s option** – The short option displays the audit event by its table name, for example, `AUE_IOCTL`. The option displays the other tokens as the default option displays them.

In the default output format of `praudit`, each record is easily identified as a sequence of audit tokens where each token is on a separate line. Each record begins with a header token. You could, for example, further process the output with the `awk` command.

Here is the default output from the `praudit` command for a header token:

```
header,240,1,ioc1(2),es,Tue Sept 7 16:11:44 1999, + 270 msec
```

Here is the output from the `praudit -r` command for the same header token:

```
20,240,1,158,0003,699754304, + 270 msec
```

Sometimes, you might want to manipulate output from the `praudit` command as lines of text. For example, you might want to select records that the `auditreduce` command cannot select. You can use a simple shell script to process the output of `praudit`. The following simple example script marks the header tokens by prefixing them with Control-A, combines all the audit tokens for one record onto one line while preserving the line breaks as Control-A, runs the `grep` command, and restores the original newline breaks.

```
#!/bin/sh
praudit | sed -e '1,2d' -e '$s/^file.*$//' -e 's/^header/^aheader/' \
| tr '\012\001' '\002\012' \
| grep "$1" \
| tr '\002' '\012'
```

Note that the `^a` in the script is Control-A, not the two characters `^` and `a`. The prefix distinguishes the header token from the string `header` that might appear as text.

## The `auditconfig` Command

The `auditconfig` command provides a command-line interface to retrieve and set audit configuration parameters. See the `auditconfig(1M)` man page. Options to the `auditconfig` command include the following:

- `-chkconf`  
Checks the configuration of kernel event-to-class mappings and reports any inconsistencies

- conf  
Reconfigures kernel event-to-class mappings at runtime to match the current mappings in the `audit_event` file.
- getcond  
Retrieves the state of auditing on the machine. The following table shows the possible responses.

**TABLE 25-1** Possible Auditing Conditions

Response	Meaning
auditing	Auditing is enabled and turned on.
no audit	Auditing is enabled, but the audit daemon is not running.
disabled	Auditing is not enabled.

- setcond *condition*  
Sets the state of auditing on the machine to one of `auditing` or `noaudit`.
- getclass *event\_number*  
Retrieves the preselection classes to which the specified event is mapped.
- setclass *event\_number audit\_flags*  
Sets the preselection classes to which the specified event is mapped.
- lsevent  
Displays the currently configured (runtime) kernel and user audit event information.
- getpinfo *pid*  
Retrieves the audit ID, preselection mask, terminal ID, and audit session ID of the specified process.
- setpmask *pid flags*  
Sets the preselection mask of the specified process.
- setsmask *asid flags*  
Sets the preselection mask of all processes with the specified audit session ID.
- setumask *auid flags*  
Sets the preselection mask of all processes with the specified user audit ID.
- lspolicy  
Displays the list of audit policies with a short description of each policy.
- getpolicy  
Shows the current audit policy flags.
- setpolicy *policy\_flag[,policy\_flag]*  
Sets the audit policy flags to the specified policies (see “Determining Which Audit Policies to Use” on page 331).

---

## Audit Files

Auditing uses the following files:

- “The `/etc/system` File” on page 363
- “The `audit_class` File” on page 363
- “The `audit_control` File” on page 364
- “The `audit_data` File” on page 365
- “The `audit_event` File” on page 365
- “The `audit_startup` Script” on page 365
- “The `audit_user` File” on page 366
- “The `audit_warn` Script” on page 367

### The `/etc/system` File

The `/etc/system` file contains commands that the kernel reads during initialization to customize the system operations. The `bsmconv` and `bsmunconv` shell scripts, which are used to activate and deactivate auditing, modify the `/etc/system` file. The `bsmconv` shell script adds the following lines to the `/etc/system` file:

```
set c2audit:audit_load=1
set abort_enable=0
```

The first command causes the `c2audit` loadable kernel module (auditing module) to be loaded when the system is booted. The second command disables the use of the `Stop-A` keyboard sequence. The `Stop-A` keyboard sequence halts the system and can invoke a debugger, which could be a security breach. The `bsmunconv` shell script removes these lines, which results in auditing being disabled when the system is rebooted.

### The `audit_class` File

The `audit_class` file contains definitions of the existing audit classes. Audit classes are groups of audit events. Each class has an associated audit flag, which is the short name that stands for the class. You use the short name in the `audit_control` file, optionally with a `+` or `-` prefix, to preselect the classes whose events you want to audit. The `root` user, or an administrator in an equivalent role, can define new audit classes, rename existing classes, or otherwise edit existing classes by using `vi`, `ed`, or some other editor. See the `audit_class(4)` man page for more information. For descriptions of the audit flags, see “Definitions of Audit Flags” on page 368.

## The audit\_control File

An `audit_control` file on each machine is read by the audit daemon (see the `audit_control(4)` man page). The `audit_control` file is located in the `/etc/security` directory. Each machine has its own local `audit_control` file so that every machine can mount their audit file systems from different locations or in a different order. For example, the primary audit file system for machineA might be the secondary audit file system for machineB.

You specify four kinds of information in the `audit_control` file:

- The *audit flags* line (`flags:`) contains the audit flags that preselect which classes of events are audited for all users on the machine. The audit flags specified here are referred to as the *machine-wide audit flags* or the *machine-wide audit preselection mask*. Audit flags are separated by commas, with no spaces.
- The *nonattributable flags* line (`naflags:`) contains the audit flags that preselect which classes of events are audited when an action cannot be attributed to a specific user. The flags are separated by commas, with no spaces.
- The *audit threshold* line (`minfree:`) defines the minimum free-space level for all audit file systems. The `minfree` percentage must be greater than or equal to 0. The default is 20 percent.
- The *directory definition* lines (`dir:`) define which audit file systems and directories the machine will use to store its audit log files. You can define one or more directory definition lines. The order of the `dir:` lines is significant. The `auditd` daemon creates audit files in the directories in the specified order (see the `audit(1M)` man page). The first directory is the primary audit directory for the machine. The second directory is the secondary audit directory where the audit daemon creates audit trail files when the first directory becomes full, and so forth.

You create an `audit_control` file during the configuration process on each machine.

When you make changes to the `audit_control` file, you then run the `audit -s` to instruct the audit daemon to reread the file.

---

**Note** – The `audit -s` command does not change the preselection mask for existing processes. Use `auditconfig`, `setaudit` (see the `getuid(2)` man page), or `auditon` for existing processes.

---

## Sample audit\_control File

The following is a sample `audit_control` file for the machine `dopey`. `dopey` uses two audit file systems on the audit server `blinken`, and a third audit file system that is mounted from the second audit server `winken`. The third file system is used only when the audit file systems on `blinken` become full or unavailable. The `minfree` value of 20 percent specifies that the warning script is run when the file systems are 80

percent filled and the audit data for the current machine will be stored in the next available audit directory, if any (see the `audit_warn(1M)` man page). The flags specify that all logins and administrative operations are to be audited (whether or not they succeed), and that failures of all types, except failures to create a file system object, are to be audited.

```
flags:lo,ad,-all,^-fc
naflags:lo,nt
minfree:20
dir:/etc/security/audit/blinken/files
dir:/etc/security/audit/blinken.1/files
#
# Audit filesystem used when blinken fills up
#
dir:/etc/security/audit/winken
```

## The `audit_data` File

When the `auditd` daemon starts on each machine, it creates the file `/etc/security/audit_data`. The format of the file consists of a single entry with the two fields separated by a colon (see the `audit_data(4)` man page). The first field is the audit daemon's process ID. The second field is the path name of the audit file to which the audit daemon is currently writing audit records. Here is an example:

```
# cat /etc/security/audit_data
116:/etc/security/audit/blinken.1/files/19990320100002.not_terminated.dopey
```

## The `audit_event` File

The `audit_event` file, located in the `/etc/security` directory, contains the default event-to-class mappings (see the `audit_event(4)` man page). You can edit this file to change the class mappings. However, if you do so, you must reboot the system or run `auditconfig -conf` to read the changed mappings into the kernel.

## The `audit_startup` Script

Auditing is enabled by starting the audit daemon, `auditd`. You can start the audit daemon by executing `/usr/sbin/auditd` as root or in an equivalent role. See the `auditd(1M)` man page.

The existence of the file `/etc/security/audit_startup` causes the audit daemon to be run automatically when the system enters multiuser mode. This file is an executable script that is invoked as part of the startup sequence, just prior to the execution of the audit daemon (see the `audit_startup(1M)` man page). A default `audit_startup` script that automatically configures the event-to-class mappings and sets the audit policies is created during the BSM package installation.

## The `audit_user` File

To audit some users differently from others, you can edit the `audit_user` file to add audit flags for individual users. If specified, these flags are combined with the system-wide flags in the `audit_control` file to determine which classes of events to audit for that user. The flags that you add to the user's entry in the `audit_user` file modify the defaults from the `audit_control` file in two ways:

- By specifying a set of event classes that are never to be audited for this user
- By specifying a set of event classes that are always to be audited for this user

Each user entry in the `audit_user` file contains three fields.

- The *username* field
- The *always-audit* field
- The *never-audit* field

The audit fields are processed in sequence. Auditing is enabled by the *always-audit* field and turned off by the *never-audit* field.

---

**Note** – Avoid the common mistake of leaving the `all` audit flag in the *never-audit* field. This mistake causes all auditing to be turned off for that user, which overrides the flags that are set in the *always-audit* field.

---

Using the *never-audit* flags for a user is not the same as removing classes from the *always-audit* set. For example, suppose you want to audit everything for user `tamiko` for except for successful reads of file system objects. This strategy audits almost everything for a user, while generating only about three-quarters of the audit data that would be produced if all data reads were audited. You also want to apply the system defaults to `tamiko`. Here are two possible `audit_user` entries:

The correct entry:

```
tamiko:all,^+fr:
```

The incorrect entry:

```
tamiko:all:+fr
```

The first example means, "always audit everything except for successful file-reads." The second example means, "always audit everything, but never audit successful file-reads." The second example is incorrect because it overrides the system default. The first example achieves the desired effect: the system defaults apply, as well as what is specified in the `audit_user` entry.

---

**Note** – Successful events and failed events are treated separately, so a process can generate more audit records when an error occurs than when an event is successful.

---

## The `audit_warn` Script

Whenever the audit daemon encounters an unusual condition while writing audit records, it invokes the `/etc/security/audit_warn` script. See the `audit_warn(1M)` man page. You can customize this script for your site to warn of conditions that might require manual intervention, or you could specify how to handle them automatically. For all error conditions, `audit_warn` writes a message to the console and sends a message to the `audit_warn` mail alias. You should set up this alias when you enable auditing.

When audit daemon detects the following conditions, it invokes the `audit_warn` script.

- An audit directory has become more full than the `minfree` value allows. The `minfree` value or soft limit is a percentage of the space available on an audit file system.

The `audit_warn` script is invoked with the string `soft` and the name of the directory whose available space has gone below the minimum value. The audit daemon switches automatically to the next suitable directory and writes the audit files there until this new directory reaches its `minfree` limit. The audit daemon then goes to each remaining directory in the order that is listed in the `audit_control` file, and writes audit records until each directory is at its `minfree` limit.

- All the audit directories have reached the `minfree` threshold.

The `audit_warn` script is invoked with the string `allsoft`. A message is written to the console and mail is sent to the `audit_warn` alias.

When all audit directories that are listed in the `audit_control` file have reached their `minfree` threshold, the audit daemon switches back to the first directory, and writes audit records until the directory becomes completely full.

- An audit directory has become completely full with no space remaining.

The `audit_warn` script is invoked with the string `hard` and the name of the directory. A message is written to the console and mail is sent to the `audit_warn` alias.

The audit daemon switches automatically to the next suitable directory with any space available, if any. The audit daemon goes to each remaining directory in the order that is listed in the `audit_control` file, and writes audit records until each directory is full.

- All the audit directories are completely full. The `audit_warn` script is invoked with the string `allhard` as an argument.

In the default configuration, a message is written to the console and mail is sent to the `audit_warn` alias. The processes that generate audit records are suspended. The audit daemon goes into a loop, waiting for space to become available, and resumes processing audit records when that happens. While audit records are not being processed, no auditable activities occur. Every process that attempts to generate an audit record is suspended. For this reason, you should set up a separate audit administration account that could operate without any auditing enabled. Then, you could continue operations without being suspended.

- An internal error occurs, such as:
  - Another audit daemon process is already running (string `ebusy`)
  - A temporary file cannot be used (string `tmpfile`)
  - The `auditsvc()` system call fails (string `auditsvc`)
  - A signal was received during auditing shutdown (string `postsigterm`)

Mail is sent to the `audit_warn` alias.

- A problem is discovered with the syntax of the `audit_control` file. By default, mail is sent to the `audit_warn` alias and a message is sent to the console.

---

## Audit Flags

*Audit flags* indicate classes of events to audit. Machine-wide defaults for auditing are specified for all users on each machine by flags in the `audit_control` file, which is described in “The `audit_control` File” on page 364.

You can modify what is audited for individual users by putting audit flags in a user’s entry in the `audit_user` file. The audit flags are also used as arguments to the `auditconfig` command (see the `auditconfig(1M)` man page).

## Definitions of Audit Flags

The following table shows each predefined audit class with the audit flag (which is the short name that stands for the class), the long name, and a short description. You use these audit flags in the auditing configuration files to specify which classes of events to audit. You can define new classes and rename existing classes by modifying the `audit_class` file (see the `audit_class(4)` man page).



**TABLE 25-2** Audit Flags

Short Name	Long Name	Short Description
no	no_class	Null value for turning off event preselection
fr	file_read	Read of data, open for reading
fw	file_write	Write of data, open for writing
fa	file_attr_acc	Access of object attributes: stat, pathconf
fm	file_attr_mod	Change of object attributes: chown, flock
fc	file_creation	Creation of object
fd	file_deletion	Deletion of object
cl	file_close	close system call
pc	process	Process operations: fork, exec, exit
nt	network	Network events: bind, connect, accept
ip	ipc	System V IPC operations
na	non_attrib	Nonattributable events
ad	administrative	Administrative actions
lo	login_logout	Login and logout events
ap	application	Application-defined event
io	ioctl	ioctl system call
ex	exec	Program execution
ot	other	Miscellaneous
all	all	All flags set

## Audit Flag Syntax

The prefixes determine whether a class of events is audited whether it succeeds or fails, or only if it succeeds, or only if it fails. Here is the format of the audit flag:

*prefixflag*

The following table shows prefixes that specify whether the audit class is audited for success or failure, or both.

**TABLE 25-3** Prefixes Used in Audit Flags

Prefix	Definition
none	Audit for both success and failure
+	Audit for success only
-	Audit for failure only

For example, the audit flag `lo` (without any prefix) means that auditing should occur for “all successful attempts to log in and log out, and all failed attempts to log in.” You cannot fail an attempt to log out. As another example, the `-all` flag means that auditing should occur only for all failed attempts of any kind. The `+all` flag means that auditing should only occur for all successful attempts of any kind.



**Caution** – The `-all` flag can generate large amounts of data and fill up audit file systems quickly. Use the `-all` flag only if you have extraordinary reasons to audit all activities.

## Prefixes to Modify Audit Flags

Use the following prefixes in any of three ways:

- In the `flags` line in the `audit_control` file to modify already specified flags
- In `flags` field in the user’s entry in the `audit_user` file
- With arguments to the `auditconfig` command

See the `auditconfig(1M)` man page.

The prefixes in the following table, along with the short names of audit classes, turn on or turn off previously specified audit classes.

**TABLE 25-4** Prefixes Used to Modify Already-Specified Audit Flags

Prefix	Definition
<code>^-</code>	Turn off for failed attempts
<code>^+</code>	Turn off for successful attempts
<code>^</code>	Turn off for both failed and successful attempts

The `^-` prefix is used in the `flags` line in the following example from an `audit_control` file.

In the following example, the `lo` and `ad` flags specify that all logins and administrative operations are to be audited when they succeed and when they fail. The `-all` means audit “all failed events.” Because the `^-` prefix means “turn off auditing for the specified class for failed attempts,” the `^-fc` flag modifies the previous flag that specified the auditing of all failed events. The two fields together mean: “audit all failed events, except for failed attempts to create file system objects.”

```
flags:lo,ad,-all,^-fc
```

---

## Audit Policies

Audit policies determine if additional tokens or information is added to the audit trail. Each policy is defined in “Determining Which Audit Policies to Use” on page 331.

---

## Process Audit Characteristics

The following audit characteristics are set at initial login:

- **Process preselection mask** – When a user logs in, the `login` command combines the machine-wide audit flags from the `audit_control` file with the user-specific audit flags (if any) from the `audit_user` file, to establish the *process preselection mask* for the user’s processes. The process preselection mask specifies whether events in each audit event class are to generate audit records.

The algorithm for obtaining the process preselection mask is as follows: the audit flags from the `flags:` line in the `audit_control` file are added to the flags from the *always-audit* field in the user’s entry in the `audit_user` file. Then, the flags from the *never-audit* field from the user’s entry in the `audit_user` file are subtracted from the total, as follows:

*user’s process preselection mask = (flags: line + always audit flags) - never audit flags*

- **Audit ID** – A process also acquires its audit ID when the user logs in. This audit ID is inherited by all child processes that were started by the user’s initial process. The audit ID helps enforce accountability. Even after a user becomes root, the audit ID remains the same. The audit ID that is saved in each audit record always allows you to trace actions back to the original user who had logged in.
- **Audit Session ID** – The audit session ID is assigned at login and inherited by all descendant processes.
- **Terminal ID (port ID, machine ID)** –The terminal ID consists of the host name and the Internet address, followed by a unique number that identifies the physical device on which the user logged in. Most often, the login is through the console and the number that corresponds to the console device is 0.

---

## Audit Trail

The *audit trail* is created by the audit daemon (see the `auditd(1M)` man page). The audit daemon starts on each machine when the machine is brought up. After the `auditd` daemon starts at boot time, it is responsible for collecting the audit trail data and writing the audit records into *audit files*, which are also called *audit log files*. For a description of the file format, see the `audit.log(4)` man page.

Even though you can physically locate audit directories within other file systems that are not dedicated to auditing, do not do so except for directories of last resort. Directories of last resort are directories where audit files are written only when there is no other suitable directory available.

There is one other scenario where locating audit directories outside of dedicated audit file systems could be acceptable. You might do so in a software development environment where auditing is optional, and where it is more important to make full use of disk space than to keep an audit trail. Putting audit directories within other file systems is not acceptable in a security-conscious production environment.

You should also consider the following:

- A host should have at least one local audit directory, which it can use as a directory of last resort if the host is unable to communicate with the audit server.
- Mount audit directories with the read-write (`rw`) option. When you mount audit directories remotely (using NFS software), also use the `intr` and `noac` options.
- List the audit file systems on the audit server where they reside. The export list should include all machines in the configuration.

## More About the Audit Files

Each audit file is a self-contained collection of records. The file's name identifies the time span during which the records were generated and the machine that generated them.

### Audit File Naming

Audit files that are complete have names of the following form:

*start-time.finish-time.machine*

where *start-time* is the time that the first audit record in the audit file was generated, *finish-time* is the time that the last record was written to the file, and *machine* is the name of the machine that generated the file. For an example of these names, see "Example of a Closed Audit File Name" on page 373.

An audit log file that is still active has a name of the following form:

*start-time.not\_terminated.machine*

## How Audit File Names Are Used

The file name time stamps are used by the `auditreduce` command to locate files that contain records for the specific time range that has been requested. These time stamps are important because there can be a month's accumulation or more of audit files online, and searching them all for records that were generated in the last 24 hours would be unacceptably expensive.

## Time-Stamp Format and Interpretation

The *start-time* and *end-time* are timestamps with one-second resolution. They are specified in Greenwich Mean Time (GMT). The format is four digits for the year, followed by two digits for each month, day, hour, minute, and second, as follows:

*YYYYMMDDHHMMSS*

The timestamps are in GMT to ensure that they will sort in proper order even across a daylight savings time boundary. Because they are in GMT, the date and hour must be translated to the current time zone to be meaningful. Beware of this point whenever you manipulate these files with standard file commands rather than with the `auditreduce` command.

## Example of a File Name for a Still-Active File

The format of a file name of a still-active file is as follows:

*YYYYMMDDHHMMSS.not\_terminated.machine*

Here is an example:

*19990327225243.not\_terminated.dopey*

The audit log files are named by the beginning date. So, in the example above the audit file was created in 1999, on March 27, at 10:52:43 p.m, GMT. The `not_terminated` in the file name means either that the file is still active or that the `auditd` daemon was unexpectedly interrupted. The name `dopey` at the end is the host name of the machine whose audit data is being collected.

## Example of a Closed Audit File Name

The format of the name of a closed audit log file is as follows:

*YYYYMMDDHHMMSS.YYYYMMDDHHMMSS.hostname*

Here is an example:

19990320005243.19990327225351.dopey

In this example, the audit log file was created in 1999, on March 20, at 12:52:43 a.m., GMT. The file was closed March 27, at 10:53:51 p.m., GMT. The name `dopey` at the end is the host name of the machine whose audit data was collected.

Whenever `auditd` is unexpectedly interrupted, the audit file that is open at the time retains the `not_terminated` file name designation. Also, when a machine is writing to a remotely mounted audit file and the file server crashes or becomes inaccessible, the `not_terminated` designation remains in the current file's name. The audit daemon opens a new audit file and keeps the old audit file name intact.

---

## Audit Record Structure

An audit record is a sequence of audit tokens. Each audit token contains event information such as user ID, time, and date. A `header` token begins an audit record, and an optional `trailer` token concludes the record. Other audit tokens contain information relevant to the auditable event. The following figure shows a typical audit record.

header token
arg token
data token
subject token
return token

**FIGURE 25-3** Typical Audit Record Structure

---

## Audit Token Formats

Each token has a token type identifier followed by data that is specific to the token. Each token type has its own format. The following table shows the token names with a description of each token.

**TABLE 25-5** Audit Tokens for the Basic Security Module

<b>Token Name</b>	<b>Description</b>	<b>For More Information</b>
acl	Access Control List information	"acl Token" on page 376
arbitrary	Data with format and type information	"arbitrary Token" on page 376
arg	System call argument value	"arg Token" on page 377
attr	file vnode tokens	"attr Token" on page 378
exec_args	Exec system call arguments	"exec_args Token" on page 378
exec_env	Exec system call environment variables	"exec_env Token" on page 379
exit	Program exit information	"exit Token" on page 380
file	Audit file information	"file Token" on page 380
groups	Process groups information (now obsolete)	"groups Token (Obsolete)" on page 381
header	Indicates start of audit record	"header Token" on page 381
in_addr	Internet address	"in_addr Token" on page 382
ip	IP header information	"ip Token" on page 382
ipc	System V IPC information	"ipc Token" on page 383
ipc_perm	System V IPC object tokens	"ipc_perm Token" on page 384
ipport	Internet port address	"ipport Token" on page 384
newgroups	Process groups information	"newgroups Token" on page 385
opaque	Unstructured data (unspecified format)	"opaque Token" on page 385
path	Path information	"path Token" on page 386
process	Process token information	"process Token" on page 386
return	Status of system call	"return Token" on page 388
seq	Sequence number token	"seq Token" on page 388
socket	Socket type and addresses	"socket Token" on page 389
subject	Subject token information (same format as process token)	"subject Token" on page 389
text	ASCII string	"text Token" on page 390
trailer	Indicates end of audit record	"trailer Token" on page 391

An audit record always contains a `header` token. The `header` token indicates where the audit record begins in the audit trail. Every audit record contains a `subject` token, except for audit records from some nonattributable events. In the case of attributable events, these two tokens refer to the values of the process that caused the event. In the case of asynchronous events, the `process` tokens refer to the system.

## acl Token

The `acl` token records information about Access Control Lists. This token consists of four fixed fields. The fixed fields are:

- a token ID that identifies this token as an `acl` token
- a field that specifies the ACL type
- an ACL ID field
- a field that lists the permissions associated with this ACL

The `praudit` command displays the `acl` token as follows:

```
acl, tpanero, staff, 0755
```

The `acl` token format is as follows:

Token ID	ACL type	ACL ID	ACL permissions
----------	----------	--------	-----------------

**FIGURE 25-4** `acl` Token Format

## arbitrary Token

The `arbitrary` token encapsulates data for the audit trail. This token consists of four fixed fields and an array of data. The fixed fields are as follows:


- a token ID that identifies this token as an `arbitrary` token
- a suggested format field (for example, hexadecimal)
- a size field that specifies the size of the data that is encapsulated (for example, short)
- a count field that provides the number of following items

The remainder of the token is composed of one or more items of the specified type. The `praudit` command displays the `arbitrary` token as follows:

```
arbitrary, decimal, int, 1  
42
```



The format of arbitrary token is as follows:



**FIGURE 25-5** arbitrary Token Format

The following table shows the possible values of the print format field. Table 25-6.

**TABLE 25-6** Values for the arbitrary Token's Print Format Field

Value	Action
AUP_BINARY	Prints the date in binary format
AUP_OCTAL	Prints the date in octal format
AUP_DECIMAL	Prints the date in decimal format
AUP_HEX	Prints the date in hexadecimal format
AUP_STRING	Prints the date as a string

The following table shows the possible values of the item size field.

**TABLE 25-7** Values for the arbitrary Token's Item Size Field

Value	Action
AUR_BYTE	Data is printed in units of bytes (1 byte)
AUR_SHORT	Data is printed in units of shorts (2 bytes)
AUR_LONG	Data is printed in units of longs (4 bytes)

## arg Token

The `arg` token contains system call argument information: the argument number of the system call, the argument value, and an optional description. This token allows a 32-bit integer system-call argument in an audit record. The `arg` token has five fields:

- a token ID that identifies this token as an `arg` token
- an argument ID that tells which system call argument the token refers to
- the argument value
- the length of the descriptive text string
- the text string

The `praudit` command displays the `arg` token as follows:

```
argument,1,0x00000000,addr
```

The following figure shows the format of the `arg` token.

Token ID	Argument #	Argument value	Text length	Text
----------	------------	----------------	-------------	------

**FIGURE 25-6** `arg` Token Format

## `attr` Token

The `attr` token contains information from the file `vnode`. This token has seven fields:

- a token ID that identifies this token as an `attr` token
- the file access mode and type
- the owner user ID
- the owner group ID
- the file system ID
- the inode ID
- the device ID the file might represent

See the `statvfs(2)` man page for further information about the file system ID and the device ID.

The `attr` token usually accompanies a `path` token and is produced during path searches. In the event of a path-search error, the `attr` token is not included as part of the audit record since there is no `vnode` available to obtain the necessary file information. The `praudit` command displays the `attr` token as follows:

```
attribute,100555,root,staff,1805,13871,-4288
```

The following figure shows the format of an `attr` token.

Token ID	File mode	Owner UID	Owner GID	File system ID	File inode ID	Device ID
----------	-----------	-----------	-----------	----------------	---------------	-----------

**FIGURE 25-7** `attr` Token Format

## `exec_args` Token

The `exec_args` token records the arguments to an `exec()` system call. The `exec_args` token has two fixed fields:

- a token ID field that identifies this token as an `exec_args` token
- a count that represents the number of arguments that are passed to the `exec()` system call

The remainder of this token is composed of zero or more null-terminated strings. The `praudit` command displays the `exec_args` token as follows:

```
vi,/etc/security/audit_user
```

The following figure shows the format of an `exec_args` token.

Token ID	Count	env_args
----------	-------	----------

**FIGURE 25-8** `exec_args` Token Format

---

**Note** – The `exec_args` token is output only when the audit policy `argv` is active.

---

## `exec_env` Token

The `exec_env` token records the current environment variables to an `exec()` system call. The `exec_env` token has two fixed fields:

- a token ID field that identifies this token as an `exec_env` token
- a count that represents the number of arguments that are passed to the `exec()` system call

The remainder of this token is composed of zero or more null-terminated strings. The `praudit` command displays the `exec_env` token as follows:

```
exec_env,25,  
GROUP=staff,HOME=/export/home/matrix,HOST=mestrix,HOSTTYPE=sun4,HZ=100,  
LC_COLLATE=en_US.ISO8859-1,LC_CTYPE=en_US.ISO8859-1,LC_MESSAGES=C,  
LC_MONETARY=en_US.ISO8859-1,LC_NUMERIC=en_US.ISO8859-1,  
LC_TIME=en_US.ISO8859-1,LOGNAME=matrix,MACHTYPE=sparc,  
MAIL=/var/mail/matrix,OSTYPE=solaris,PATH=/usr/sbin:/usr/bin,PS1=#,  
PWD=/var/audit,REMOTEHOST=209.198.087.208,SHELL=/usr/bin/csh,SHLVL=1,  
TERM=dtterm,TZ=US/Pacific,USER=matrix,VENDOR=sun
```

The following figures shows the format of an `exec_env` token.

Token ID	Count	env_args
----------	-------	----------

**FIGURE 25-9** `exec_env` Token Format

---

**Note** – The `exec_env` token is output only when the audit policy `arge` is active.

---

## exit Token

The `exit` token records the exit status of a program. The `exit` token contains the following fields:

- a token ID that identifies this token as an `exit` token
- a program exit status as passed to the `exit()` system call
- a return value that describes the exit status or provides a system error number

The `praudit` command displays the `exit` token as follows:

```
exit,Error 0,0
```

The following figure shows the format of an `exit` token.

Token ID	Status	Return value
----------	--------	--------------

**FIGURE 25–10** `exit` Token Format

## file Token

The `file` token is a special token that is generated by the audit daemon to mark the beginning of a new audit trail file and the end of an old audit trail file as it is deactivated. The audit daemon builds a special audit record that contains this token to “link” together successive audit files into one audit trail. The `file` token has four fields:

- a token ID that identifies this token as a `file` token
- a time and date stamp that identifies the time that the file was created or closed
- a byte count of the file name that includes a null terminator
- a field that holds the file null-terminated name

The `praudit` command displays the `file` token as follows:

```
file,Tue Sep  1 13:32:42 1992, + 79249 msec,  
  /var/audit/localhost/files/19990901202558.19990901203241.quisp
```

The following figure shows the format of a `file` token.

Token ID	Date & time	Name length	Previous/next file name
----------	-------------	-------------	-------------------------

**FIGURE 25–11** `file` Token Format

## groups Token (Obsolete)

This token has been replaced by the `newgroups` token, which provides the same type of information but requires less space. A description of the `groups` token is provided here for completeness, but the application designer should use the `newgroups` token. Notice that `praudit` does not distinguish between the two tokens, as both token IDs are labelled `groups` when ASCII output is displayed.

The `groups` token records the `groups` entries from the process's credential. The `groups` token has two fixed fields:

- A token ID that identifies this token as a `groups` token
- An array of `groups` entries of size `NGROUPS_MAX` (16)

The remainder of the token consists of zero or more group entries. The `praudit` command displays the `group` token as follows:

```
group, staff, admin, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1, -1
```

The following figure shows the format of a `groups` token.

Token ID	Groups
----------	--------

**FIGURE 25-12** `groups` Token Format

---

**Note** – The `groups` token is output only when the audit policy group is active.

---

## header Token

The `header` token is special in that it marks the beginning of an audit record and combines with the `trailer` token to bracket all the other tokens in the record. The `header` token has six fields:

- a token ID field that identifies this token as a `header` token
- a byte count of the total length of the audit record, including both the header and the trailer
- a version number that identifies the version of the audit record structure
- the audit event ID that identifies the type of audit event that the record represents
- the ID modifier that identifies special characteristics of the audit event
- and the time and date that the record was created

On 64-bit systems, the `header` token is displayed with a 64-bit time stamp, in place of the 32-bit time stamp.

The `praudit` command displays the header token for a `ioctl()` system call as follows:

```
header,240,1,ioctl(2),es,Tue Sept 1 16:11:44 2001, + 270000 msec
```

The following figure shows the format of a header token.

Token ID	Byte count	Version #	Event ID	ID modifier	Date and time
----------	------------	-----------	----------	-------------	---------------

**FIGURE 25-13** header Token Format

The ID modifier field has the following flags defined:

0x4000	PAD_NOTATTR	nonattributable event
0x8000	PAD_FAILURE	fail audit event

## in\_addr Token

The `in_addr` token contains a 4-byte Internet Protocol address. The `in_addr` token has two fields:

- a token ID that identifies this token as an `ip` address token
- an Internet address

The `praudit` command displays the `in_addr` token as follows:

```
ip address,129.150.113.7
```

For the Solaris 8 release, the Internet address can be displayed as an IPv4 address that uses 4 bytes, or as an IPv6 address that uses 16 bytes to describe the type, and 16 bytes to describe the address. The following figure shows the format of an `in_addr` token.

Token ID	Internet address
----------	------------------

**FIGURE 25-14** `in_addr` Token Format

## ip Token


The `ip` token contains a copy of an Internet Protocol header but does not include any IP options. The IP options can be added by including more of the IP header in the token. The `ip` token has two fields:

- a token ID that identifies this token as an `ip` token
- a copy of the IP header (all 20 bytes)

The `praudit` command displays the `ip` token as follows:

ip address,0.0.0.0

The IP header structure is defined in the `/usr/include/netinet/ip.h` file. The following figure shows the format of an ip token.



**FIGURE 25-15** ip Token Format

## ipc Token

The `ipc` token contains the System V IPC message/semaphore/shared-memory handle that is used by the caller to identify a particular IPC object. The `ipc` token has three fields:

- a token ID that identifies this token as an IPC token
- a type field that specifies the type of IPC object
- the handle that identifies the IPC object

The `praudit` command displays the `ipc` token as follows:

```
IPC,msg,3
```

---

**Note** – The IPC object identifiers violate the context-free nature of the Solaris CMW audit tokens. No global “name” uniquely identifies IPC objects. Instead, they are identified by their handles, which are valid only during the time the IPC objects are active. The identification should not be a problem since the System V IPC mechanisms are seldom used, and they all share the same audit class.

---

The following table shows the possible values for the IPC object type field. The values are defined in the `/usr/include/bsm/audit.h` file.

**TABLE 25-8** Values for the IPC Object Type Field

Name	Value	Description
AU_IPC_MSG	1	IPC message object
AU_IPC_SEM	2	IPC semaphore object
AU_IPC_SHM	3	IPC shared-memory object

The following figure shows the format of an `ipc` token.

Token ID	IPC object type	IPC object ID
----------	-----------------	---------------

**FIGURE 25-16** `ipc` Token Format

## `ipc_perm` Token

The `ipc_perm` token contains a copy of the System V IPC access information. This token is added to audit records that are generated by IPC shared-memory events, IPC semaphore events, and IPC message events. The `ipc_perm` token has eight fields:

- a token ID that identifies this token as an `ipc_perm` token
- the user ID of the IPC owner
- the group ID of the IPC owner
- the user ID of the IPC creator
- the group ID of the IPC creator
- the access modes of the IPC
- the sequence number of the IPC
- the IPC key value

The `praudit` command displays the `ipc_perm` token as follows:

```
IPC perm,root,wheel,root,wheel,0,0,0x00000000
```

The values are taken from the `ipc_perm` structure that is associated with the IPC object. The following figure shows the format of an `ipc_perm` token.

Token ID	Owner uid	Owner gid	Creator uid	Creator gid	IPC mode	Sequence ID	IPC key
----------	-----------	-----------	-------------	-------------	----------	-------------	---------

**FIGURE 25-17** `ipc_perm` Token Format

## `ipport` Token

The `ipport` token contains the TCP (or UDP) port address. The `ipport` token has two fields:

- a token ID that identifies this token as an `ipport` token
- the TCP/UDP port address

The `praudit` command displays the `ipport` token as follows:

```
ip port,0xf6d6
```



The following figure shows the format of an `ipport` token.

Token ID	Port ID
----------	---------

**FIGURE 25-18** `ipport` Token Format

## newgroups Token

This token replaces the `groups` token. Notice that the `praudit` command does not distinguish between the two tokens, as both token IDs are labelled `groups` when ASCII output is displayed.

The `newgroups` token records the `groups` entries from the process's credential. The `newgroups` token has two fixed fields:

- a token ID field that identifies this token as a `newgroups` token
- a count that represents the number of groups that are contained in this audit record

The remainder of this token is composed of zero or more group entries. The `praudit` command displays the `ipport` token as follows:

```
group, staff, admin
```

The following figure shows the format of a `newgroups` token.

Token ID	Count	Groups
----------	-------	--------

**FIGURE 25-19** `newgroups` Token Format

---

**Note** – The `newgroups` token is output only when the `group` audit policy is active.

---

## opaque Token


The `opaque` token contains unformatted data as a sequence of bytes. The `opaque` token has three fields:

- a token ID that identifies this token as an `opaque` token
- a byte count of the data
- an array of byte data

The `praudit` command displays the `opaque` token as follows:

```
opaque, 12, 0x4f5041515545204441544100
```

The following figure shows the format of an opaque token.



**FIGURE 25-20** opaque Token Format

## path Token


The path token contains access path information for an object. This token contains the following fields:

- a token ID that identifies this token as an path token
- a byte count of the path length
- the absolute path to the object that is based on the real root of the system

The `praudit` command displays the path token as follows. Note that the path length field is not displayed.

```
path,/etc/security/audit_user
```

The following figure shows the format of a path token.



**FIGURE 25-21** path Token Format

## process Token

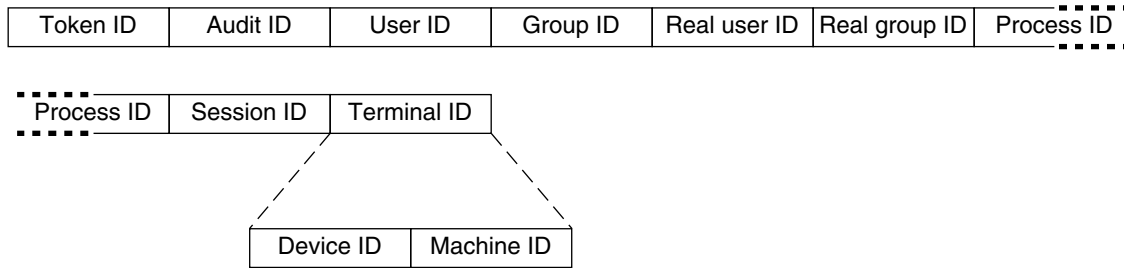
The process token contains information about a user who is associated with a process, such as the recipient of a signal. The process token has nine fields:

- a token ID that identifies this token as a process token
- the invariant audit ID
- the effective user ID
- the effective group ID
- the real user ID
- the real group ID
- the process ID
- the audit session ID
- a terminal ID that consists of a device ID and a machine ID

The `praudit` command displays the `process` token as follows:

```
process,root,root,wheel,root,wheel,0,0,0,0.0.0.0
```

The following figure shows the format of a `process` token.



**FIGURE 25–22** `process` Token Format

The audit ID, user ID, group ID, process ID, and session ID are long instead of short.

---

**Note** – The `process` token fields for the session ID, the real user ID, or the real group ID might be unavailable. The value is then set to -1.

---

Any token that contains a terminal ID has several variations. The `praudit` command hides these variations on output of the terminal ID so that they all appear the same. This field is handled the same way for any token that contains it. The terminal ID is either an IP address and port number, or a device ID, such as the serial port that is connected to a modem, in which case it is zero. The terminal ID is specified in one of several formats:

For device numbers:

- 32-bit applications: 4-byte device number, 4-bytes unused
- 64-bit applications: 8-byte device number, 4-bytes unused

For port numbers in the Solaris 7 release or earlier releases:

- 32-bit applications: 4-byte port number, 4-byte IP address
- 64-bit applications: 8-byte port number, 4-byte IP address

For port numbers in the Solaris 8 or 9 releases:

- 32-bit with IPV4: 4-byte port number, 4-byte IP type, 4-byte IP address
- 32-bit with IPV6: 4-byte port number, 4-byte IP type, 16-byte IP address
- 64-bit with IPV4: 8-byte port number, 4-byte IP type, 4-byte IP address
- 64-bit with IPV6: 8-byte port number, 4-byte IP type, 16-byte IP address

## return Token

The return token contains the return status of the system call (`u_error`) and the process return value (`u_rval1`). This token has three fields:

- a token ID that identifies this token as a return token
- the error status of the system call
- the system call return value

The return token is always returned as part of kernel-generated audit records for system calls. This token indicates exit status and other return values in application auditing.

The `praudit` command displays the return token as follows:

```
return, success, 0
```

The following figure shows the format of a return token.

Token ID	Process error	Process value
----------	---------------	---------------

**FIGURE 25-23** return Token Format

## seq Token

The seq token (sequence token) is an optional token that contains a sequence number. Used for debugging, this token is added to each audit record when the seq policy is active. The seq token has two fields:

- a token ID that identifies this token as a seq token
- a 32-bit unsigned long field that contains the sequence number

The sequence number is incremented every time an audit record is generated and added to the audit trail. The `praudit` command displays the seq token as follows:

```
sequence, 1292
```

The following figure shows the format of a seq token.

Token ID	Sequence number
----------	-----------------

**FIGURE 25-24** seq Token Format

---

**Note** – The seq token is output only when the seq audit policy is active.

---

## socket Token

The `socket` token contains information that describes an Internet socket. This token has six fields:

- a token ID that identifies this token as a `socket` token
- a socket type field that indicates the type of socket referenced (TCP/UDP/UNIX)
- the local port address
- the local Internet address
- the remote port address
- the remote Internet address

The `praudit` command displays the `socket` token as follows:

```
socket,0x0000,0x0000,0.0.0.0,0x0000,0.0.0.0
```

For the Solaris 8 release, the Internet address can be displayed as a IPv4 address that uses 4 bytes, or as an IPv6 address that uses 16 bytes to describe the type, and 16 bytes to describe the addresses. The following figure shows the format of a `socket` token.

Token ID	Type	Remote port	Remote Internet address
----------	------	-------------	-------------------------

**FIGURE 25–25** `socket` Token Format

## subject Token

The `subject` token describes a user who performs or attempts to perform an operation. The format is the same as the `process` token. The `subject` token has nine fields:

- an ID that identifies this token as a `subject` token
- the invariant audit ID
- the effective user ID
- the effective group ID
- the real user ID
- the real group ID
- the process ID
- the audit session ID
- a terminal ID that consists of a device ID and a machine ID

The `subject` token is always returned as part of kernel-generated audit records for system calls. The `praudit` command displays the `subject` token as follows:

```
subject,cjc,cjc,staff,cjc,staff,424,223,0 0 quisp
```

The audit ID, user ID, group ID, process ID, and session ID are long instead of short.

---

**Note** – The `subject` token fields for the session ID, the real user ID, or the real group ID might be unavailable. The value is then set to -1.

---

Any token that contains a terminal ID has several variations. The `praudit` command hides these variations on output of the terminal ID so that they all appear the same. This field is handled the same way for any token that contains it. The terminal ID is either an IP address and port number, or a device ID, such as the serial port that is connected to a modem, in which case it is zero. The terminal ID is specified in one of several formats:

For device numbers:

- 32-bit applications: 4-byte device number, 4-bytes unused
- 64-bit applications: 8-byte device number, 4-bytes unused

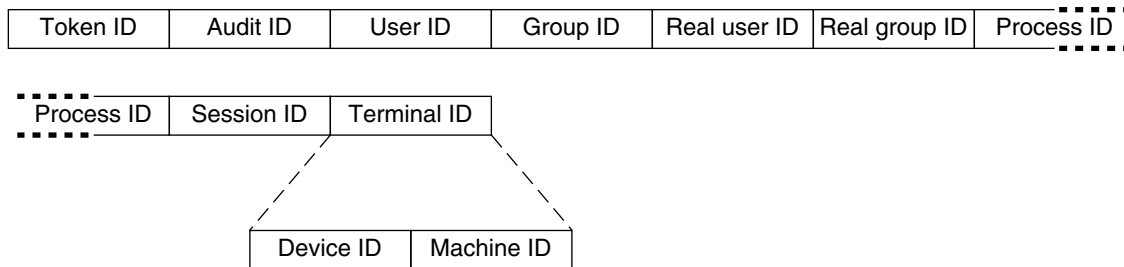
For port numbers in the Solaris 7 release or earlier releases:

- 32-bit applications: 4-byte port number, 4-byte IP address
- 64-bit applications: 8-byte port number, 4-byte IP address

For port numbers in the Solaris 8 or 9 releases:

- 32-bit with IPV4: 4-byte port number, 4-byte IP type, 4-byte IP address
- 32-bit with IPV6: 4-byte port number, 4-byte IP type, 16-byte IP address
- 64-bit with IPV4: 8-byte port number, 4-byte IP type, 4-byte IP address
- 64-bit with IPV6: 8-byte port number, 4-byte IP type, 16-byte IP address

The following figure shows the format of the `subject` token.



**FIGURE 25-26** `subject` Token Format

## text Token

The `text` token contains a text string. This token has three fields:

- a token ID that identifies this token as a `text` token

- the length of the text string
- the text string itself

The `praudit` command displays the `text` token as follows:

```
text,aw_test_token
```

The following figure shows the format of a `text` token.

Token ID	Text length	Text string
----------	-------------	-------------

**FIGURE 25-27** text Token Format

## trailer Token

The two tokens, `header` and `trailer`, are special in that they distinguish the end points of an audit record and bracket all the other tokens. A `header` token begins an audit record. A `trailer` token ends an audit record. The `trailer` token is an optional token and is added as the last token of each record only when the `trail` audit policy has been set.

The `trailer` token supports backward seeks of the audit trail. The `trailer` token has three fields:

- a token ID that identifies this token as a `trailer` token
- a pad number to aid in marking the end of the record
- the total number of characters in the audit record, including both the `header` and `trailer` tokens

The `praudit` command displays the `trailer` token as follows:

```
trailer,136
```

The following figure shows the format of a `trailer` token.

Token ID	Pad number	Byte count
----------	------------	------------

**FIGURE 25-28** trailer Token Format

The audit trail analysis software ensures that each record contains both the `header` and `trailer` tokens. In the case of a write error, as when a file system becomes full, an audit record can be incomplete and truncated. The `auditsvc()` system call, that is responsible for writing data to the audit trail, attempts to write complete audit records. When file system space runs out, the system call terminates without releasing the current audit record. When the system call resumes, it can then repeat the truncated record. For more information, see the `auditsvc(2)` man page.

---

## Device Allocation Reference

Device allocation protects removable media from unauthorized use. By requiring that a user allocate a device, or by denying a user permission to use a device, you can protect your site from loss of data, computer viruses, and other security breaches. The following section provides information about BSM device allocation.

### Components of the Device-Allocation Mechanism

The components of the device-allocation mechanism are as follows:

- The `allocate`, `deallocate`, `dminfo`, and `list_devices` commands. For more information, see “Using the Device-Allocation Commands” on page 392.
- The `/etc/security/device_allocate` file. See the `device_allocate(4)` man page.
- The `/etc/security/device_maps` file. See the `device_maps(4)` man page.
- A lock file in the `/etc/security/dev` directory for each allocatable device.
- The changed attributes of the *device-special files* that are associated with each allocatable device.
- Device-clean scripts for each allocatable device.

The `device_allocate` file, the `device_map` file, and the lock files are local configuration files. These files are not administered as name service databases because tape drives, diskette drives, and printers are all connected to specific machines.

### Using the Device-Allocation Commands

This section describes some of the options to the `allocate`, `deallocate`, and `list_devices` commands that are usable only by `root` or a role of equivalent power. The commands are detailed on their respective man pages.

`allocate -F device_special_filename`

Reallocates the specified device. This option is often used with the `-U` option to reallocate the specified device to the specified user. Without the `-U` option, the device is allocated to `root`.

`allocate -U username`

Causes the device to be allocated to the user who is specified rather than to the current user. This option allows you to allocate a device for another user while you are `root`, without having to assume that user’s identity.



`deallocate -F device_special_filename`

Forces the deallocation of a device. Devices that a user has allocated are not automatically deallocated when the process terminates or when the user logs out. When a user forgets to deallocate a tape drive, you can force deallocation by using the `-F` option while you are `root`.

`deallocate -I`

Forces the deallocation of all allocatable devices. This option should be used only at system initialization.

`list_devices`

Lists all the device-special files that are associated with any device that is listed in the `device_maps` file.

`list_devices -U username`

Lists the devices that are allocatable or allocated to the user ID that is associated with the specified user name. This option allows you to check which devices are allocatable or allocated to another user while you are `root`.

## The Allocate Error State

An allocatable device is in the *allocate error state* if it is owned by user `bin` and group `bin` with a device-special file mode of `0100`. If a user wants to allocate a device that is in the allocate error state, you should try to force the deallocation of the device, by using the `deallocate` command with the `-F` option. Or, you can use `allocate -U` to assign the device to the user, then investigate any error messages that appear. After any problems with the device are corrected, you must rerun the `deallocate -F` or `allocate -F` command to clear the allocate error state from the device.

## The device\_maps File

You can examine the `/etc/security/device_maps` file to determine the device names, device types, and device-special files that are associated with each allocatable device. See the `device_maps(4)` man page. Device maps are created when you set up device allocation. A rudimentary `device_maps` file is created by `bsmconv` when the BSM is enabled. This initial `device_maps` file should be used only as a starting point. You can then augment and customize the `device_maps` file for your site.

The `device_maps` file defines the device-special file mappings for each device, which in many cases is not intuitive. This file allows various programs to discover which device-special files map to which devices. You can use the `dminfo` command, for example, to retrieve the device name, the device type, and the device-special files to specify when you set up an allocatable device. The `dminfo` command uses the `device_maps` file to report this information.

Each device is represented by a one-line entry of the form:

*device-name:device-type:device-list*

Lines in the `device_maps` file can end with a backslash (\) to continue an entry on the next line. Comments can also be included. A “#” makes a comment out of all subsequent text until the next newline not immediately preceded by a backslash. Leading and trailing blanks are allowed in any field.

*device-name*

Specifies the name of the device, for example `st0`, `fd0`, or `audio`. The device name that is specified here must correspond to the name of the lock file that is used in the `/etc/security/dev` directory.

*device-type*

Specifies the generic device type (which is the name for the class of devices, such as `st`, `fd`, `audio`). The *device-type* field logically groups related devices.

*device-list*

Lists of the device-special files that are associated with the physical device. The *device-list* must contain *all* of the special files that allow access to a particular device. If the list is incomplete, a malevolent user can still obtain or modify private information. Valid entries for the *device-list* field are either the real device files located under `/devices` or the symbolic links that are in `/dev`. The symbolic links in the `/dev` directory are provided for binary compatibility.

The following is an example of a `device_maps` file entries for SCSI tape `st0` and diskette `fd0`.

```
fd0:\
    fd:\
    /dev/fd0 /dev/fd0a /dev/fd0b /dev/rfd0 /dev/rfd0a /dev/rfd0b:\
        .
        .
        .
st0:\
    st:\
    /dev/rst0 /dev/rst8 /dev/rst16 /dev/nrst0 /dev/nrst8 /dev/nrst16:\
```

## The `device_allocate` File

You can modify the `device_allocate` file to change devices from allocatable to nonallocatable, or to add new devices. A sample `device_allocate` file follows.

```
st0;st;;;/etc/security/lib/st_clean
fd0;fd;;;/etc/security/lib/fd_clean
sr0;sr;;;/etc/security/lib/sr_clean
audio;audio;;;*/etc/security/lib/audio_clean
```

You define which devices should be allocatable during initial BSM configuration. You can decide to accept the default devices and their defined characteristics, as shown in the preceding sample `device_allocate` file. Whenever you add a device to any machine after the system is up and running, you must decide whether to make the new device allocatable.

After installation, you can modify the entries for devices in the `device_allocate` file. Any device that needs to be allocated before use must be defined in the `device_allocate` file on each machine. Currently, cartridge tape drives, diskette drives, CD-ROM devices, and audio chips are considered allocatable and have device-clean scripts.

---

**Note** – Xylogics™ tape drives or Archive tape drives also use the `st_clean` script that is supplied for SCSI devices. Other devices that you can make allocatable are modems, terminals, graphics tablets, and similar devices. However, you need to create your own device-clean scripts for such devices, and the script must fulfill object-reuse requirements for that type of device.

---

An entry in the `device_allocate` file does not mean that the device is allocatable, unless the entry specifically states that the device is allocatable. In the sample `device_allocate` file, note the asterisk (\*) in the fifth field of the audio device entry. An asterisk in the fifth field indicates to the system that the device is not allocatable. That is, the system administrator does not require a user to allocate the device before it is used nor to deallocate it afterward. Any other string placed in this field indicates that the device is allocatable.

In the `device_allocate` file, you represent each device by a one-line entry of the form:

*device-name ; device-type ; reserved ; reserved ; alloc ; device-clean*

For example, the following line shows the entry for device name `st0`:

```
st0;st;;;;;/etc/security/lib/st_clean
```

Lines in the `device_allocate` file can end with a “\” to continue an entry on the next line. Comments can also be included. A “#” makes a comment out of all subsequent text until the next newline not immediately preceded by a “\”. Leading and trailing blanks are allowed in any field.

The following paragraphs describe each field in the `device_allocate` file.

*device-name*

Specifies the name of the device; for example, `st0`, `fd0`, or `sr0`. When you make a device allocatable, retrieve the *device-name* from the *device-name* field in the `device_maps` file, or use the `dminfo` command. (The name is also the DAC file name for the device.)

#### *device-type*

Specifies the generic device type (which is the name for the class of devices, such as `st`, `fd`, and `sr`). This field groups related devices. When you make an allocatable device, retrieve the *device-type* from the *device-type* field in the `device_maps` file, or use the `dmINFO` command.

#### *reserved*

These two fields are reserved for future use.

#### *alloc*

Specifies whether the device is allocatable. An asterisk (\*) in this field indicates that the device is *not* allocatable. Any other string, or an empty field, indicates that the device is allocatable.

#### *device-clean*

Supplies the path name of a script to be invoked for special handling, such as cleanup and object-reuse protection during the allocation process. The *device-clean* script is run any time that the device is acted on by the `deallocate` command, such as when a device is forcibly deallocated with `deallocate -F`.

## Device-Clean Scripts

The *device-clean* scripts address the security requirement that all usable data be purged from a physical device before reuse. By default, cartridge tape drives, diskette drives, CD-ROM devices, and audio devices require device-clean scripts, which are provided. This section describes what device-clean scripts do.

## Object Reuse

Device allocation satisfies part of the object-reuse requirement. The device-clean scripts make sure that data that is left on a device by one user is cleared before the device is allocatable by another user.

### *Device-Clean Script for Tapes*

The `st_clean` device-clean script supports three tape devices. The supported tape devices are:

- SCSI 1/4-inch tape
- Archive 1/4-inch tape
- Open-reel 1/2-inch tape

The `st_clean` script uses the `rewoffl` option to the `mt` command to affect the device cleanup. For more information, see the `mt(1)` man page. If the script runs during system boot, it queries the device to see if the device is online and has media in it. The 1/4-inch tape devices that have media in them, are placed in the allocate error state to force the administrator to clean up the device manually.

During normal system operation, when the `allocate` or `deallocate` command is executed in interactive mode, the user is prompted to remove the media from the device that is being deallocated. The script pauses until the media is removed from the device.

### *Device-Clean Scripts for Diskettes and CD-ROM Devices*

The following table shows the device-clean scripts for diskettes and CD-ROM devices.

**TABLE 25-9** Device-Clean Scripts for Diskettes and CD-ROM Devices

Disk Device Type	Device-Clean Script
Diskette	<code>fd_clean</code>
CD-ROM	<code>sr_clean</code>

The scripts use the `eject` command to remove the media from the drive. See the `eject(1)` man page. If the `eject` command fails, the device is placed in the `allocate` error state.

### *Device-Clean Script for Audio*

Audio devices are cleaned up with an `audio-clean` script. The script performs an `AUDIO_DRAIN` ioctl system call to flush the device, then an `AUDIO_SETINFO` ioctl system call command to reset the device configuration to default. In addition, the script retrieves the audio chip registers by using the `AUDIOGETREG` ioctl system call. Any registers that deviate from the default settings are reset by using the `AUDIOSETREG` ioctl system call.

## Writing New Device-Clean Scripts

If you add more allocatable devices to the system, you might need to create your own device-clean scripts. The `deallocate` command passes a parameter to the device-clean scripts. The parameter, shown here, is a string that contains the device name (see the `device_allocate(4)` man page):

```
st_clean -[I|F|S] device-name
```

Device-clean scripts must return "0" for success and greater than "0" for failure. The options `-I`, `-F`, and `-S` help the script determine its running mode.

The `-I` option is needed during system boot only. All output must go to the system console. Failure or inability to forcibly eject the media must put the device in the `allocate` error state.

The `-F` option is for forced cleanup. This option is interactive and assumes that the user is available to respond to prompts. A script with this option must attempt to complete the cleanup if one part of the cleanup fails.

The `-S` option is for standard cleanup. This option is interactive and assumes that the user is available to respond to prompts.

## How the Device-Allocation Mechanism Works

This section gives an example of how the device-allocate mechanism works.

The `allocate` command first checks for the presence of a lock file under the device name for the specified device in the `/etc/security/dev` directory. If the file is owned by `allocate`, then the ownership of the lock file is changed to the name of the user who initiated the `allocate` command.

The `allocate` command then checks for an entry for the device in the `device_allocate` file, and checks whether the entry shows that the device as allocatable.

The first listing in the following example shows that a lock file exists with owner `bin`, group `bin`, and mode `600` for the `st0` device in `/etc/security/dev`. The second listing shows that the associated device-special files are set up properly, with owner `bin`, group `bin`, and mode `000`.

```
untouchable% ls -lg /etc/security/dev/st0
-rw----- 1 bin bin                0 Dec 6 15:21 /etc/security/dev/st0
untouchable% ls -lg /devices/sbus@1,f8000000/esp@0,800000
c----- 1 bin bin                18,  4 May 12 13:11 st@4,0:
c----- 1 bin bin                18, 20 May 12 13:11 st@4,0:b
c----- 1 bin bin                18, 28 May 12 13:11 st@4,0:bn
c----- 1 bin bin                18, 12 May 12 13:11 st@4,0:c
.
.
.
c----- 1 bin bin                18,  0 May 12 13:11 st@4,0:u
c----- 1 bin bin                18, 16 May 12 13:11 st@4,0:ub
c----- 1 bin bin                18, 24 May 12 13:11 st@4,0:ubn
c----- 1 bin bin                18,  8 May 12 13:11 st@4,0:un
```

In this example, the user `vanessa` allocates device `st0`.

```
untouchable% whoami
vanessa
untouchable% allocate st0
```

When the user `vanessa` runs the `allocate` command to allocate the tape `st0`, `allocate` first checks for the existence of an `/etc/security/dev/st0` file. If no lock file exists or if the lock file is owned by a user other than `allocate`, then user `vanessa` could not allocate the device.

If the `allocate` command finds the lock file for the device with the correct ownership and permissions, it then checks to make sure that the device has an entry in the `device_allocate` file and that the entry specifies that the device is allocatable.

In this example, the default `device_allocate` entry for the `st0` device specifies that the device is allocatable. Because the `allocate` command finds that all these conditions are met, the device is allocated to user `vanessa`.

The `allocate` command changes the ownership and permissions of the device-special files that are associated with the device in the `/dev` directory. To allocate the `st0` device to the user `vanessa`, the mode on its associated device-special files is changed to `600` and the owner is changed to `vanessa`.

The `allocate` command also changes the ownership of the lock file that is associated with the device in the `/etc/security/dev` directory. To allocate the `st0` device to the user `vanessa`, the owner of `/etc/security/dev/st0` is changed to `vanessa`.

In the following example, after the user `vanessa` executes the `allocate` command with the device name `st0`, the owner of `/etc/security/dev/st0` is changed to `vanessa` and the owner of the associated device-special files is now also `vanessa`. Lastly, user `vanessa` now has permission to read and write the files.

```
untouchable% whoami
vanessa
untouchable% allocate st0
untouchable% ls -lg /etc/security/dev/st0
-rw----- 1 vanessa staff          0 Dec 6 15:21 /etc/security/dev/st0
untouchable% ls -la /devices/sbus@1,f8000000/esp@0,800000
.
.
.
crw----- 1 vanessa 18,  4 May 12 13:11 st@4,0:
crw----- 1 vanessa 18, 12 May 12 13:11 st@4,0:b
crw----- 1 vanessa 18, 12 May 12 13:11 st@4,0:bn
crw----- 1 vanessa 18, 12 May 12 13:11 st@4,0:c
.
.
.
crw----- 1 vanessa 18,  4 May 12 13:11 st@4,0:u
crw----- 1 vanessa 18, 12 May 12 13:11 st@4,0:ub
crw----- 1 vanessa 18, 12 May 12 13:11 st@4,0:ubn
crw----- 1 vanessa 18, 12 May 12 13:11 st@4,0:un
```





# Glossary

---

<b>admin principal</b>	A user principal with a name of the form <i>username/admin</i> (as in <i>joe/admin</i> ). An admin principal can have more privileges (for example, to change policies) than a regular user principal. See also principal name, user principal.
<b>application server</b>	See network application server.
<b>authentication</b>	The process of verifying the claimed identity of a principal.
<b>authenticator</b>	Authenticators are passed by clients when requesting tickets (from a KDC) and services (from a server). They contain information that is generated by using a session key known only by the client and server, that can be verified as of recent origin, thus indicating that the transaction is secure. When used with a ticket, an authenticator can be used to authenticate a user principal. An authenticator includes the principal name of the user, the IP address of the user's host, and a timestamp. Unlike a ticket, an authenticator can be used only once, usually when access to a service is requested. An authenticator is encrypted by using the session key for that client and that server.
<b>authorization</b>	<ol style="list-style-type: none"><li>1. The process of determining if a principal can use a service, which objects the principal is allowed to access, and the type of access that is allowed for each object.</li><li>2. In role-based access control (RBAC), a permission that can be assigned to a role or user (or embedded in a rights profile) for performing a class of actions that are otherwise prohibited by security policy.</li></ol>
<b>client</b>	<p>Narrowly, a process that makes use of a network service on behalf of a user; for example, an application that uses <i>rlogin</i>. In some cases, a server can itself be a client of some other server or service.</p> <p>More broadly, a host that a) receives a Kerberos credential, and b) makes use of a service that is provided by a server.</p>

	Informally, a principal that makes use of a service.
<b>client principal</b>	(RPCSEC_GSS API) A client (a user or an application) that uses RPCSEC_GSS-secured network services. Client principal names are stored in the form of <code>rpc_gss_principal_t</code> structures.
<b>clock skew</b>	The maximum amount of time that the internal system clocks on all hosts that are participating in the Kerberos authentication system can differ. If the clock skew is exceeded between any of the participating hosts, requests are rejected. Clock skew can be specified in the <code>krb5.conf</code> file.
<b>confidentiality</b>	See privacy.
<b>credential</b>	An information package that includes a ticket and a matching session key. Used to authenticate the identity of a principal. See also ticket, session key.
<b>credential cache</b>	A storage space (usually a file) that contains credentials that are received from the KDC.
<b>flavor</b>	Historically, <i>security flavor</i> and <i>authentication flavor</i> had the same meaning, as a flavor that indicated a type of authentication (AUTH_UNIX, AUTH_DES, AUTH_KERB). RPCSEC_GSS is also a security flavor, even though it provides integrity and privacy services in addition to authentication.
<b>forwardable ticket</b>	A ticket that a client can use to request a ticket on a remote host without requiring the client to go through the full authentication process on that host. For example, if the user <code>david</code> obtains a forwardable ticket while on user <code>jennifer</code> 's machine, he can log in to his own machine without being required to get a new ticket (and thus authenticate himself again). See also proxiable ticket.
<b>FQDN</b>	Fully qualified domain name. For example, <code>denver.mtn.example.com</code> (as opposed to simply <code>denver</code> ).
<b>GSS-API</b>	The Generic Security Service Application Programming Interface. A network layer that provides support for various modular security services (including SEAM). GSS-API provides for security authentication, integrity, and privacy services. See also authentication, integrity, privacy.
<b>host</b>	A machine that is accessible over a network.
<b>host principal</b>	A particular instance of a service principal in which the principal (signified by the primary name <code>host</code> ) is set up to provide a range of network services, such as <code>ftp</code> , <code>rcp</code> , or <code>rlogin</code> . An example of a host principal is <code>host/boston.eng.example.com@ENG.EXAMPLE.COM</code> . See also server principal.
<b>initial ticket</b>	A ticket that is issued directly (that is, not based on an existing ticket-granting ticket). Some services, such as applications that change

passwords, might require tickets to be marked *initial* so as to assure themselves that the client can demonstrate a knowledge of its secret key. This assurance is important because an initial ticket indicates that the client has recently authenticated itself (instead of relying on a ticket-granting ticket, which might exist for a long time).

<b>instance</b>	The second part of a principal name, an instance qualifies the principal's primary. In the case of a service principal, the instance is required and is the host's fully qualified domain name, as in <code>host/boston.eng.example.com</code> . For user principals, an instance is optional. Note, however, that <code>joe</code> and <code>joe/admin</code> are unique principals. See also primary, principal name, service principal, user principal.
<b>integrity</b>	A security service that, in addition to user authentication, provides for the validity of transmitted data through cryptographic checksumming. See also authentication, privacy.
<b>invalid ticket</b>	A postdated ticket that has not yet become usable. An invalid ticket is rejected by an application server until it becomes validated. To be validated, an invalid ticket must be presented to the KDC by the client in a TGS request, with the <code>VALIDATE</code> flag set, after its start time has passed. See also postdated ticket.
<b>KDC</b>	Key Distribution Center. A machine that has three Kerberos V5 components: <ul style="list-style-type: none"><li>■ Principal and key database</li><li>■ Authentication service</li><li>■ Ticket-granting service</li></ul> Each realm has a master KDC and should have one or more slave KDCs.
<b>Kerberos</b>	An authentication service, the protocol that is used by that service, or the code that is used to implement that service.  SEAM is an authentication implementation that is closely based on Kerberos V5.  While technically different, "SEAM" and "Kerberos" are often used interchangeably in SEAM documentation. The same is true for "Kerberos" and "Kerberos V5."  Kerberos (also spelled Cerberus) was a fierce, three-headed mastiff who guarded the gates of Hades in Greek mythology.
<b>key</b>	<ol style="list-style-type: none"><li>1. An entry (principal name) in a keytab file. See also keytab file.</li><li>2. An encryption key, of which there are three types:</li></ol>

	<ul style="list-style-type: none"> <li>■ A <i>private key</i> – An encryption key that is shared by a principal and the KDC, and distributed outside the bounds of the system. See also private key.</li> <li>■ A <i>service key</i> – This key serves the same purpose as the private key, but is used by servers and services. See also service key.</li> <li>■ A <i>session key</i> – A temporary encryption key that is used between two principals, with a lifetime limited to the duration of a single login session. See also session key.</li> </ul>
<b>keytab file</b>	A key table file that contains one or more keys (principals). A host or service uses a keytab file in the much the same way that a user uses a password.
<b>kvno</b>	Key version number. A sequence number that tracks a particular key in order of generation. The highest kvno is the latest and most current key.
<b>name service scope</b>	The scope in which a role is permitted to operate, that is, an individual host or all hosts that are served by a specified name service such as NIS, NIS+, or LDAP. Scopes are applied to Solaris Management Console toolboxes.
<b>master KDC</b>	The main KDC in each realm, which includes a Kerberos administration server, <code>kadmin</code> , and an authentication and ticket-granting daemon, <code>krb5kdc</code> . Each realm must have at least one master KDC, and can have many duplicate, or slave, KDCs that provide authentication services to clients.
<b>mechanism</b>	A software package that specifies cryptographic techniques to achieve data authentication or confidentiality. Examples: Kerberos V5, Diffie-Hellman public key.
<b>network application server</b>	A server that provides a network application, such as <code>ftp</code> . A realm can contain several network application servers.
<b>NTP</b>	Network Time Protocol. Software from the University of Delaware that enables you to manage precise time or network clock synchronization, or both, in a network environment. You can use NTP to maintain clock skew in a Kerberos environment. See also clock skew.
<b>PAM</b>	Pluggable Authentication Module. A framework that allows for multiple authentication mechanisms to be used without having to recompile the services that use them. PAM enables SEAM session initialization at login.
<b>policy</b>	A set of rules, initiated when SEAM is installed or administered, that govern ticket usage. Policies can regulate principals' accesses, or ticket parameters, such as lifetime.
<b>postdated ticket</b>	A postdated ticket does not become valid until some specified time after its creation. Such a ticket is useful, for example, for batch jobs that

are intended to run late at night, since the ticket, if stolen, cannot be used until the batch job is run. When a postdated ticket is issued, it is issued as *invalid* and remains that way until a) its start time has passed, and b) the client requests validation by the KDC. A postdated ticket is normally valid until the expiration time of the ticket-granting ticket. However, if the postdated ticket is marked *renewable*, its lifetime is normally set to be equal to the duration of the full life time of the ticket-granting ticket. See also *invalid ticket*, *renewable ticket*.

<b>primary</b>	The first part of a principal name. See also <i>instance</i> , <i>principal name</i> , <i>realm</i> .
<b>principal</b>	<ol style="list-style-type: none"><li>1. A uniquely named client/user or server/service instance that participates in a network communication. Kerberos transactions involve interactions between principals (service principals and user principals) or between principals and KDCs. In other words, a principal is a unique entity to which Kerberos can assign tickets. See also <i>principal name</i>, <i>service principal</i>, <i>user principal</i>.</li><li>2. (RPCSEC_GSS API) See <i>client principal</i>, <i>server principal</i>.</li></ol>
<b>principal name</b>	<ol style="list-style-type: none"><li>1. The name of a principal, in the format <i>primary/instance@REALM</i>. See also <i>instance</i>, <i>primary</i>, <i>realm</i>.</li><li>2. (RPCSEC_GSS API) See <i>client principal</i>, <i>server principal</i>.</li></ol>
<b>privacy</b>	A security service, in which transmitted data is encrypted before being sent. Privacy also includes data integrity and user authentication. See also <i>authentication</i> , <i>integrity</i> , <i>service</i> .
<b>private key</b>	A key that is given to each user principal, and known only to the user of the principal and to the KDC. For user principals, the key is based on the user's password. See also <i>key</i> .
<b>private-key encryption</b>	In private-key encryption, the sender and receiver use the same key for encryption. See also <i>public-key encryption</i> .
<b>privileged application</b>	An application that can override system controls and that checks for specific UIDs, GIDs, or authorizations.
<b>profile shell</b>	A shell used in RBAC that enables a role (or user) to run any privileged applications that are assigned to the role's rights profiles from the command line. The profile shells are <i>pfsh</i> , <i>pfssh</i> , and <i>pfksh</i> . They correspond to the Bourne shell ( <i>sh</i> ), C shell ( <i>csh</i> ), and Korn shell ( <i>ksh</i> ), respectively.
<b>proxiable ticket</b>	A ticket that can be used by a service on behalf of a client to perform an operation for the client. Thus, the service is said to act as the client's proxy. With the ticket, the service can take on the identity of the client. The service can use a proxiable ticket to obtain a service ticket to another service, but it cannot obtain a ticket-granting ticket. The

difference between a proxiable ticket and a forwardable ticket is that a proxiable ticket is only valid for a single operation. See also forwardable ticket.

<b>public-key encryption</b>	An encryption scheme in which each user has two keys, one public key and one private key. In public-key encryption, the sender uses the receiver's public key to encrypt the message, and the receiver uses a private key to decrypt it. SEAM is a private-key system. See also private-key encryption.
<b>QOP</b>	Quality of Protection. A parameter that is used to select the cryptographic algorithms that are used in conjunction with the integrity service or privacy service.
<b>RBAC</b>	Role-Based Access Control. An alternative to the all-or-nothing superuser model. RBAC lets an organization separate superuser's capabilities and assign them to special user accounts called roles. Roles can be assigned to specific individuals, according to their job needs.
<b>realm</b>	<ol style="list-style-type: none"><li>1. The logical network that is served by a single SEAM database and a set of Key Distribution Centers (KDCs).</li><li>2. The third part of a principal name. For the principal name <code>joe/admin@ENG.EXAMPLE.COM</code>, the realm is <code>ENG.EXAMPLE.COM</code>. See also principal name.</li></ol>
<b>relation</b>	A configuration variable or relationship that is defined in the <code>kdc.conf</code> or <code>krb5.conf</code> files.
<b>renewable ticket</b>	Because having tickets with very long lives is a security risk, tickets can be designated as <i>renewable</i> . A renewable ticket has two expiration times: a) the time at which the current instance of the ticket expires, and b) maximum lifetime for any ticket. If a client wants to continue to use a ticket, the client renews the ticket before the first expiration occurs. For example, a ticket can be valid for one hour, with all tickets having a maximum lifetime of ten hours. If the client that holds the ticket wants to keep it for more than an hour, the client must renew the ticket. When a ticket reaches the maximum ticket lifetime, it automatically expires and cannot be renewed.
<b>rights profile</b>	Also referred to as right or profile. A collection of overrides used in RBAC that can be assigned to a role or user. A rights profile can consist of authorizations, commands with set UIDs or GIDs, which are referred to as security attributes, and other rights profiles.
<b>role</b>	A special identity for running privileged applications that only assigned users can assume.

<b>SEAM</b>	<p>Sun Enterprise Authentication Mechanism. A system for authenticating users over a network, based on the Kerberos V5 technology that was developed at the Massachusetts Institute of Technology.</p> <p>“SEAM” and “Kerberos” are often used interchangeably in the SEAM documentation.</p>
<b>secret key</b>	See private key.
<b>Secure Shell</b>	A special protocol for secure remote login and other secure network services over an insecure network.
<b>security flavor</b>	See flavor.
<b>security mechanism</b>	See mechanism.
<b>security service</b>	See service.
<b>server</b>	A principal that provides a resource to network clients. For example, if you <code>rlogin</code> to the machine <code>boston.eng.acme.com</code> , then that machine is the server that provides the <code>rlogin</code> service. See also service principal.
<b>server principal</b>	(RPCSEC_GSS API) A principal that provides a service. The server principal is stored as an ASCII string in the form <code>service@host</code> . See also client principal.
<b>service</b>	<ol style="list-style-type: none"> <li>1. A resource that is provided to network clients, often by more than one server. For example, if you <code>rlogin</code> to the machine <code>boston.eng.example.com</code>, then that machine is the server that provides the <code>rlogin</code> service.</li> <li>2. A security service (either integrity or privacy) that provides a level of protection beyond authentication. See also integrity and privacy.</li> </ol>
<b>service key</b>	An encryption key that is shared by a service principal and the KDC, and is distributed outside the bounds of the system. See also key.
<b>service principal</b>	A principal that provides Kerberos authentication for a service or services. For service principals, the primary name is a name of a service, such as <code>ftp</code> , and its instance is the fully qualified host name of the system that provides the service. See also host principal, user principal.
<b>session key</b>	A key that is generated by the authentication service or the ticket-granting service. A session key is generated to provide secure transactions between a client and a service. The lifetime of a session key is limited to a single login session. See also key.
<b>slave KDC</b>	A copy of a master KDC, which is capable of performing most functions of the master. Each realm usually has several slave KDCs (and only one master KDC). See also KDC, master KDC.

<b>stash file</b>	A stash file contains an encrypted copy of the master key for the KDC. This master key is used when a server is rebooted to automatically authenticate the KDC before it starts the <code>kadmind</code> and <code>krb5kdc</code> processes. Because the stash file includes the master key, the stash file and any backups of it should be kept secure. If the encryption is compromised, then the key could be used to access or modify the KDC database.
<b>ticket</b>	An information packet that is used to securely pass the identity of a user to a server or service. A ticket is valid for only a single client and a particular service on a specific server. A ticket contains the principal name of the service, the principal name of the user, the IP address of the user's host, a timestamp, and a value that defines the lifetime of the ticket. A ticket is created with a random session key to be used by the client and the service. Once a ticket has been created, it can be reused until the ticket expires. A ticket only serves to authenticate a client when it is presented along with a fresh authenticator. See also authenticator, credential, service, session key.
<b>ticket file</b>	See credential cache.
<b>TGS</b>	Ticket-Granting Service. That portion of the KDC that is responsible for issuing tickets.
<b>TGT</b>	Ticket-Granting Ticket. A ticket that is issued by the KDC that enables a client to request tickets for other services.
<b>user principal</b>	A principal that is attributed to a particular user. A user principal's primary name is a user name, and its optional instance is a name that is used to describe the intended use of the corresponding credentials (for example, <code>joe</code> or <code>joe/admin</code> ). Also known as a user instance. See also service principal.
<b>VPN</b>	Virtual Private Network. A network that provides secure communication by using encryption and tunneling to connect users over a public network.



# Index

---

## Numbers and Symbols

- \* (asterisk)
  - device\_allocate file, 395
  - wildcard character in ASET, 312
- ^+ audit flag prefix, 370
- # (pound sign)
  - device\_allocate file, 395
  - device\_maps file, 394
- ? (question mark)
  - in ASET tune files, 312
- + audit flag prefix, 370
- \ (backslash)
  - device\_allocate file, 395
  - ending in device\_maps file, 394
- . (dot)
  - path variable entry, 199
- ^- audit flag prefix, 370
- ~/ .gkadmin file
  - description, 181
- \$HOME/ .ssh/known\_hosts file
  - description, 68
- ~/ .k5login file
  - description, 181
- 3des-cbc encryption algorithm
  - ssh\_config file, 64
- 3des encryption algorithm
  - sshd\_config file, 64

## A

- absolute mode
  - changing file permissions, 219, 221

- absolute mode (*continued*)
  - description, 219
  - setting special permissions, 219
- access
  - getting to server
    - with SEAM, 190
  - obtaining for a specific service, 192
  - restricting for KDC servers, 120
  - root access
    - displaying attempts on console, 243
    - monitoring su command use, 205, 243
    - restricting, 209, 241
- security
  - ACLs, 201, 227
  - file access restriction, 198
  - firewall setup, 200
  - login access restrictions, 202
  - login control, 198
  - monitoring system usage, 199
  - network control, 198
  - path variable setting, 199
  - physical site security, 198
  - reporting problems, 200
  - root access restrictions, 241
  - root login tracking, 205
  - setuid programs, 199
- sharing files, 209
- system logins, 203
- access control list, *See* ACL
- Access Control Lists (ACLs), *See* ACL
- ACL
  - adding entries, 232
  - changing entries, 232

- ACL (*continued*)
  - checking entries, 231
  - commands, 202
  - default entries for directories, 228
  - deleting entries, 202, 233
  - description, 201, 227
  - directory entries, 228
  - displaying entries, 202, 233
  - format of entries, 227
  - kadm5.acl file, 139, 146, 148, 152
  - setting entries, 229, 231
  - valid file entries, 227
- acl token
  - format, 376
- ad audit flag, 369
- Add Administrative Role wizard
  - description, 259, 265
- Add Right dialog box
  - description, 271
- Add User wizard
  - description, 257
- adding
  - administration principals (SEAM), 93
  - allocatable devices (BSM), 352
  - custom roles (RBAC), 267
  - PAM module, 40
  - rights profiles (RBAC), 273
  - roles (RBAC), 265, 267
  - service principal to keytab file (SEAM), 168
  - the first role (RBAC), 259
  - the first user (RBAC), 257
- admin\_server section
  - krb5.conf file, 92
- administering
  - BSM
    - audit class, 326
    - audit classes, 368
    - audit event, 326
    - audit files, 349
    - audit flags, 327, 368
    - audit records, 328
    - audit trail overflow prevention, 352
    - auditreduce command, 349
    - cost control, 333
    - description, 324
    - efficiency, 334
    - kernel events, 327
    - process preselection mask, 334
  - administering, BSM (*continued*)
    - reducing storage-space requirements, 334
    - user-level events, 327
  - SEAM
    - keytabs, 166
    - policies, 153
    - principals, 140
  - Secure Shell, 61
  - administrative audit class, 369
  - aes128-cbc encryption algorithm
    - ssh\_config file, 64
  - agent daemon
    - Secure Shell, 54
  - aliases file (ASET)
    - description, 303
    - example, 313
    - format, 313
    - specification, 306
  - all
    - audit class, 369
    - audit flag
      - caution for using, 370
      - described, 369
      - in user audit fields, 366
  - All rights profile
    - description, 280
  - allhard string
    - audit\_warn script, 367
  - allocate command
    - authorizations required, 294
    - how the allocate mechanism works, 398
    - options, 392
    - using, 354
  - allocate error state, 393
  - AllowGroups keyword
    - sshd\_config file, 66
  - AllowTCPForwarding keyword
    - sshd\_config file, 66
  - AllowUsers keyword
    - sshd\_config file, 66
  - allsoft string
    - audit\_warn script, 367
  - always-audit flags
    - description, 366
    - process preselection mask, 371
  - analysis
    - praudit command, 360

- ap audit flag, 369
- application audit class, 369
- arbitrary token
  - format, 376
  - item size field, 377
  - print format field, 377
- Archive tape drive clean script, 395
- arg token, 378
- arge audit policy
  - description, 332
  - exec\_env token and, 379
- argv audit policy
  - description, 332
  - exec\_args token and, 379
- ASET
  - description, 295
  - environment variables, 308
  - error messages, 316
  - NFS servers and, 307
- aset command
  - initiating ASET sessions, 296
  - p option, 315
  - running ASET interactively, 313
  - running ASET periodically, 314
  - stop running ASET periodically, 315
- aset.restore command
  - description, 307
- ASETDIR variable (ASET)
  - working directory specification, 309
- asetenv file
  - description, 304
  - modifying, 304
  - running ASET periodically, 314
- ASETSECLEVEL variable (ASET)
  - setting security levels, 309
- Assign Administrative Role dialog box
  - description, 268
- Assign Rights to Role dialog box
  - description, 269
- asterisk (\*)
  - device\_allocate file, 395
  - wildcard character, 312
- at command
  - authorizations required, 294
- atq command
  - authorizations required, 294
- attr token, 378
- audio\_clean script, 397
- audio devices
  - device-clean scripts, 397
  - AUDIO\_DRAIN ioctl system call, 397
  - AUDIO\_SETINFO ioctl system call, 397
  - AUDIOGETREG ioctl system call, 397
  - AUDIOSETREG ioctl system call, 397
  - audit flag prefix, 370
- audit characteristics
  - overview, 371
  - process preselection mask, 334
- audit class
  - description, 326
- audit classes
  - auditconfig command options, 362
  - description, 325
  - flags and definitions, 368
  - mapping events, 327
- audit command
  - n option, 358
  - preselection mask for existing processes (-s option), 364
  - rereading audit files (-s option), 358
  - resetting directory pointer (-s option), 358
- audit\_control file
  - audit daemon rereading after editing, 364
  - audit\_user file modification, 366
- dir: line
  - described, 364
  - examples, 365
- examples, 365
- flags: line
  - described, 364
  - prefixes in, 370
  - process preselection mask, 371
- minfree: line
  - audit\_warn condition, 367
  - described, 364
- naflags: line, 364
- overview, 324, 364
- prefixes in flags line, 370
- problem with contents, 368
- audit daemon
  - audit\_startup file, 365
  - audit trail creation, 357, 372
  - audit\_warn script
    - conditions invoking, 367
    - described, 367
    - execution of, 358

- audit daemon (*continued*)
  - enabling auditing, 365
  - functions, 357
  - order audit files are opened, 364
  - rereading the `audit_control` file, 364
- `audit_data` file, 365
- audit directory
  - description, 325
- audit event
  - `audit_event` file, 326
  - description, 325
  - kernel event, 327
  - mapping to classes, 327
  - user-level events, 327
- `audit_event` file, 326
- audit events
  - kernel events
    - `auditconfig` command options, 361
    - `auditconfig` command options, 362
  - user-level events
    - `auditconfig` command options, 362
- audit files
  - `auditreduce` command, 358, 360
  - combining, 350, 358, 360
  - copying messages to single file, 350
  - displaying in entirety, 349
  - file token, 380
  - minimum free space for file systems, 364
  - names, 372
    - form, 373
    - still-active files, 373
  - nonactive files marked `not_terminated`, 351
  - order for opening, 364
  - printing, 349
  - reducing, 350, 358, 360
  - reducing storage-space requirements, 334
  - switching to new file, 358
  - time stamps, 373
- audit flags, 368
  - `audit_control` file line, 364
  - `audit_user` file, 366
  - `auditconfig` command options, 362
  - definitions, 368
  - description, 325
  - machine-wide, 327, 364, 368
  - overview, 327, 368
  - prefixes, 370
  - process preselection mask, 371
- audit flags (*continued*)
  - syntax, 369
- audit ID, 371
  - overview, 323
- audit messages
  - copying to single file, 350
- audit policies
  - `auditconfig` options, 362
  - default, 331
  - description, 326
  - list of, 332
- audit records
  - audit directories full, 358, 367, 391
  - converting to readable format, 328, 359
  - description, 326
  - events that generate, 324
  - format or structure, 374
  - overview, 328
  - reducing audit files, 350
- audit session ID, 371
- `audit_startup` file, 365
- audit threshold, 364
- audit tokens
  - audit record format, 374
  - description, 326, 328
  - format, 374
  - table of, 374
- audit trail
  - analysis
    - `praudit` command, 360
  - analysis costs, 333
  - creating
    - audit daemon's role, 357, 372
    - `audit_data` file, 365
    - overview, 372
  - description, 326
  - events included, 327
  - merging all files, 359
  - monitoring in real time, 334
  - overflow prevention, 352
  - overview, 324
- `audit_user` file
  - prefixes for flags, 370
  - process preselection mask, 371
  - user audit fields, 366
- `audit_warn` script, 368
  - audit daemon execution of, 358
  - conditions invoking, 367

- audit\_warn script (*continued*)
  - description, 367
  - strings, 367
- auditconfig command
  - audit flags as arguments, 328, 368
  - options, 361
  - prefixes for flags, 370
- auditd daemon
  - audit\_startup file, 365
  - audit trail creation, 357, 372
  - audit\_warn script
    - conditions invoking, 367
    - described, 367
    - execution of, 358
  - enabling auditing, 365
  - functions, 357
  - order audit files are opened, 364
  - rereading the audit\_control file, 364
- auditreduce command, 358, 360
  - c option, 350
  - cleaning not\_terminated files, 351
  - d option, 350
  - description, 358
  - examples, 349
  - O option, 350
  - options, 359
  - time stamp use, 373
  - without options, 359
- auditsvc() system call
  - audit\_warn script and, 368
  - trailer token and, 391
- AUE... names
  - description, 327
- auth\_attr database
  - description, 285, 288
  - RBAC relationships, 287
- AUTH\_DH authentication, 36
- AUTH\_DH client-server session, 29, 32
  - additional transaction, 32
  - client authenticates server, 31
  - contacting the server, 30
  - decrypting the conversation key, 30
  - generating public and secret keys, 29
  - generating the conversation key, 30
  - running keylogin, 29
  - storing information on the server, 31
  - verifier returned to client, 31
- authentication
  - configuring cross-realm, 98
  - description, 208
  - DH, 28, 36
  - network security, 208
  - overview of Kerberos, 189
  - root for NFS, 109
  - SEAM and, 71
  - Secure Shell
    - description, 50
    - hosts, 50
    - methods, 51
    - steps, 62
    - users, 50
  - terminology, 185
  - types, 208
- authentication parameters
  - ssh\_config file, 63
- authenticator
  - in SEAM, 186, 191
- authorization
  - database
    - See* auth\_attr database
  - delegating, 285
  - description, 208, 246, 250, 284
  - granularity, 285
  - naming convention, 284
  - network security, 208
  - SEAM and, 71
  - types, 208
- authorized\_keys file
  - description, 68
- auths command
  - description, 292
- authtok\_check module
  - description, 42
- authtok\_get module
  - description, 42
- authtok\_store module
  - description, 42
- Automated Security Enhancement Tool, *See* ASET
- automatically enabling auditing, 365
- automating principal creation, 141

## B

- backing up the Kerberos database, 114
- backslash (\)
  - device\_allocate file, 395
- backup
  - Kerberos database, 114
  - slave KDCs, 86
- Basic Solaris User rights profile
  - description, 280, 283
- Batchmode keyword
  - ssh\_config file, 65
- blowfish-cbc encryption algorithm
  - ssh\_config file, 64
- blowfish encryption algorithm
  - ssh\_config file, 64
- Bourne shell
  - ASET working directory specification, 309
  - privileged version, 249
- bsmconv script
  - devicemaps file creation, 393

## C

- c option
  - auditreduce command, 350
- C shell
  - ASET working directory specification, 309
  - privileged version, 249
- cache
  - credential, 190
- caret (^) in audit flag prefixes, 370
- CD-ROM drives
  - device-clean scripts, 397
- cd subcommand
  - sftp command, 58
- changepw principal, 167
- changing
  - (command line) user properties, 275
  - rights profiles (command line), 274
  - role properties (command line), 270
  - your password with kpasswd, 178
  - your password with passwd, 178
- CheckHostIP keyword
  - ssh\_config file, 64
- chgrp command
  - description, 201
  - syntax, 218
- chgrp subcommand
  - sftp command, 58
- chkconf option
  - auditconfig command, 361
- chkey command, 29, 35
- chmod command
  - changing special permissions, 222
  - description, 201
  - syntax, 222
- chmod subcommand
  - sftp command, 58
- choosing
  - your password, 177
- chown command
  - description, 201
  - syntax, 217
- Cipher keyword
  - ssh\_config file, 64
- Ciphers keyword
  - ssh\_config file, 64
  - sshd\_config file, 66
- cklist.rpt file
  - description, 298, 302
- CKLISTPATH\_level variable (ASET)
  - setting the directories to be checked, 311
- cl audit flag, 369
- class
  - description, 326
- classes
  - auditconfig command options, 362
  - flags and definitions, 368
- cleaning
  - not\_terminated files, 351
- client
  - AUTH\_DH client-server session, 29, 32
  - definition in SEAM, 185
- client names
  - planning for in SEAM, 85
- clients (SEAM)
  - configuring, 106
- clock skew
  - SEAM and, 87, 109
- clock synchronizing
  - SEAM and, 87, 95, 98
- cnt audit policy
  - description, 332
- combining audit files, 350
  - auditreduce command, 358, 360

- command-line equivalents of SEAM
  - Administration Tool, 137
- commands
  - device-allocation commands, 393
  - table of SEAM, 183
- comments
  - device\_allocate file, 395
  - device\_maps file, 394
- common key
  - calculation, 31
  - DH authentication and, 28
- Compression keyword
  - ssh\_config file, 65
- CompressionLevel keyword
  - ssh\_config file, 65
- Computer Emergency Response Team/Coordination Center (CERT/CC), 200
- conf option
  - auditconfig command, 362
- configuration decisions
  - SEAM
    - client and service principal names, 85
    - clock synchronization, 87
    - database propagation, 87
    - mapping hostnames onto realms, 85
    - number of realms, 84
    - ports, 86
    - realm hierarchy, 84
    - realm names, 84
    - realms, 84
    - slave KDCs, 86
- configuration file
  - PAM, 44, 47
- configuring
  - ASET, 304, 307
  - audit trail overflow prevention, 352
  - auditconfig command, 361
  - RBAC
    - task map, 254
  - SEAM
    - adding administration principals, 93
    - clients, 106
    - cross-realm authentication, 98
    - master KDC server, 91
    - NFS servers, 101
    - overview, 89
    - slave KDC server, 95
  - configuring, SEAM (*continued*)
    - task map, 89
    - Secure Shell, 63
  - ConnectionAttempts keyword
    - ssh\_config file, 65
  - console
    - displaying su command use on, 243
    - root access restriction to, 241
  - context-sensitive help, 138
  - control flags
    - PAM, 46
  - conversation key
    - decrypting, 30
    - generating, 30
  - converting
    - audit records to readable format, 328, 359, 361
  - copying audit messages to single file, 350
  - cost control
    - BSM and, 333
  - creating
    - credential table, 103
    - /etc/d\_passwd file, 240
    - keytab file, 93
    - new policy, 146
    - new policy (SEAM), 158
    - new principal (SEAM), 146
    - Secure Shell keys, 52
    - stash file, 97
    - tickets with kinit, 174
  - creating the audit trail, 372
    - audit daemon's role, 357
    - audit\_data file, 365
    - auditd daemon, 358
    - overview, 372
  - cred database, 33
    - DH authentication and, 28
  - cred table
    - information stored by server, 31
  - credential
    - cache, 190
    - description, 30, 185
    - obtaining for a server, 191
    - obtaining for a TGS, 190
    - or tickets, 73
  - credential cache, 190
  - credential table
    - adding single entry to, 103

- cron command
  - backing up using, 97
- cron service name
  - PAM, 45
- crontab files
  - authorizations required, 294
  - running ASET periodically, 296
  - stop running ASET periodically, 315
- cross-realm authentication
  - configuring, 98
- crypt command
  - file security, 201
- csh command
  - dial-up passwords, 240
  - privileged version, 249
- .cshrc file
  - path variable entry, 199

**D**

- d option
  - auditreduce command, 350
  - praudit command, 361
- d\_passwd file
  - creating, 240
  - description, 239
  - disabling dial-up logins temporarily, 241
  - /etc/passwd file and, 239
- daemon
  - keyserv, 32
  - krb5kdc, 92
  - table of SEAM, 184
- Data Encryption Standard, *See* DES
- data forwarding
  - Secure Shell, 62
- database
  - backing up and propagating KDC, 97, 114
  - creating KDC, 92
  - KDC propagation, 87
- deallocate command
  - allocate error state, 393
  - authorizations required, 294
  - description, 393
  - device-clean scripts and, 397
  - using, 354
- debugging sequence number, 388

- decrypting
  - conversation key, 30
  - secret key, 29
- default\_realm section
  - krb5.conf file, 92
- defaults
  - ACL entries for directories, 228
  - audit\_startup file, 365
  - machine-wide, 368
  - praudit output format, 361
- delete\_entry command, 172
- deleting
  - ACL entries, 202, 233
  - host's service, 172
  - policies (SEAM), 161
  - principal (SEAM), 149
- DenyGroups keyword
  - sshd\_config file, 66
- DenyUsers keyword
  - sshd\_config file, 66
- DES encryption, 28
- destroying
  - tickets with kdestroy, 176
- device\_allocate file
  - format, 395
  - overview, 394
- device allocation, 354
  - adding devices, 352
  - allocatable devices, 352, 395
  - allocate command
    - how the allocate mechanism works, 398
    - options, 392
    - using, 354
  - allocate error state, 393
  - allocating a device, 354
  - commands, 392
  - components of the allocation
    - mechanism, 392
  - deallocate command
    - allocate error state, 393
  - deallocate command
    - allocate error state, 393
  - deallocate command
    - described, 393
  - deallocate command
    - device-clean scripts and, 397
  - deallocate command
    - using, 354



- device allocation (*continued*)
  - description, 326
  - device\_allocate file, 394
  - device-clean scripts, 398
    - audio devices, 397
    - CD-ROM drives, 397
    - described, 396
    - diskette drives, 397
    - options, 397
    - tape drives, 395
    - writing new scripts, 397
  - device\_maps file, 393
  - list\_devices command, 393
  - lock file setup, 353
  - managing devices, 353
  - reallocating, 392
  - using device allocations, 354
- device-clean scripts
  - audio devices, 397
  - CD-ROM drives, 397
  - description, 396
  - diskette drives, 397
  - options, 397
  - tape drives, 395
  - writing new scripts, 397
- device\_maps file
  - format, 393
  - overview, 393
- devices
  - device allocation
    - See* device allocation
  - lock files, 353
  - managing, 353
  - system device access control, 198, 241
- dfstab file, 104
  - kerberos option, 104
  - sharing files, 209
- DH authentication, 28
  - AUTH\_DH client-server session, 29, 32
  - mounting files, 36
  - sharing files, 35
- DH security
  - for an NIS+ client, 33
  - for an NIS client, 34
- dhkeys module
  - description, 43
- dial\_auth module
  - description, 43
- dial-up passwords, 239
  - disabling dial-up logins temporarily, 241
  - /etc/d\_passwd file, 239
  - /etc/dialups file, 239
- dialups file
  - creating, 240
  - description, 239
- Diffie-Hellman
  - role in authentication, 62
- dir: line
  - audit\_control file, 364
- direct realms, 99
- directories
  - audit\_control file definitions, 364
  - audit daemon pointer, 358
  - audit directories full, 358, 367
  - mounting audit directories, 372
- directory
  - ACL entries, 228
  - ASET files, 296
    - checklist task (CKLISTPATH)
      - setting, 305, 311
    - master files, 303
    - reports, 302
    - working directory, 309, 313
  - audit directories full, 391
  - displaying files and related information, 201, 215
  - permissions
    - defaults, 214
    - description, 212
    - public directories, 214
- disabling
  - dial-up logins temporarily, 241
  - service on a host (SEAM), 171
  - user logins, 237
- disk-space requirements, 334
- diskette drives
  - device-clean scripts, 397
- displaying
  - ACL entries, 202, 233
  - ASET task status, 297, 300
  - audit log in entirety, 349
  - files and related information, 201, 215
  - root access attempts on console, 243
  - su command use on console, 243
  - sublist of principals (SEAM), 143
  - user's login status, 235

- dminfo command, 393
- DNS, 86
  - SEAM and, 85
- domain\_realm section
  - krb5.conf file, 85, 92
- dot (.)
  - path variable entry, 199
- DSAAuthentication keyword
  - sshd\_config file, 65
- dtlogin service name
  - PAM, 45
- .dtprofile script
  - use in Secure Shell, 55
- dtsession service name
  - PAM, 46
- duplicating
  - principal (SEAM), 148

## E

- ebusy string
  - audit\_warn script, 368
- editing rights profiles
  - task description, 270
- eeprom.rpt file
  - description, 299, 302
- efficiency
  - BSM and, 334
- eject command
  - BSM device cleanup and, 397
- enabling
  - auditing, 365
- encrypting
  - capturing encrypted passwords, 241
  - files, 201
- encryption, 28
  - privacy service, 71
  - specifying algorithms in ssh\_config, 64
  - specifying algorithms in sshd\_config, 66
- ending
  - signal received during auditing
    - shutdown, 368
- env.rpt file
  - description, 299, 302
- environment file
  - description, 68

- environment file (ASET)
  - description, 304
  - modifying, 304
  - running ASET periodically, 314
- environment variables
  - ASET
    - ASETDIR, 309
    - ASETSECLEVEL, 309
    - CKLISTPATH\_level, 305, 311
    - PERIODIC\_SCHEDULE, 306, 310, 314
    - summary table, 308
    - TASKS, 305, 310
    - UID\_ALIASES, 303, 306, 311
    - YPCHECK, 306, 311
- equals sign (=)
  - file permissions symbol, 220
- error message
  - with kpasswd, 178
- errors
  - allocate error state, 393
  - audit directories full, 358, 367, 391
  - internal errors, 368
- EscapeChar keyword
  - ssh\_config file, 65
- /etc/d\_passwd file, 239
  - creating, 240
  - disabling dial-up logins temporarily, 241
  - /etc/passwd file and, 239
- /etc/default/login file
  - restricting root access to console, 242
  - restricting root access to devices, 241
- /etc/default/su file
  - displaying su command use on
    - console, 243
  - monitoring su command, 243
- /etc/dfs/dfstab file
  - kerberos option, 104
  - sharing files, 209
- /etc/dialups file, 239
  - creating, 240
- /etc/group file
  - ASET checks, 298
- /etc/hosts.equiv file
  - description, 68
- /etc/init.d/kdc file
  - description, 181
- /etc/init.d/kdc.master file
  - description, 181

- /etc/krb5/kadm5.acl file
  - description, 182
- /etc/krb5/kadm5.keytab file
  - description, 182
- /etc/krb5/kdc.conf file
  - description, 182
- /etc/krb5/kpropd.acl file
  - description, 182
- /etc/krb5/krb5.conf file
  - description, 182
- /etc/krb5/krb5.keytab file
  - description, 182
- /etc/krb5/warn.conf file
  - description, 182
- /etc/logindevperm file
  - description, 198
- /etc/nologin file, 237
  - description, 68
- /etc/nsswitch.conf file
  - login access restrictions, 202
- /etc/pam.conf
  - description, 47, 182
  - syntax, 44
- /etc/passwd file
  - ASET checks, 298
  - /etc/d\_passwd file and, 239
- /etc/publickey file
  - DH authentication and, 29
- /etc/security/audit/bsmconv script
  - devicemaps file creation, 393
- /etc/security/audit\_data file, 365
- /etc/security/audit\_event file, 326
  - audit events and, 326
- /etc/security/audit\_startup file, 365
- /etc/security/audit\_warn script, 367
- /etc/security/dev lock files, 353
- /etc/ssh\_host\_key.pub file
  - description, 67
- /etc/ssh/shosts.equiv file
  - description, 68
- /etc/ssh/ssh\_config file
  - client authentication parameters, 63
  - configuring Secure Shell, 63
  - host-specific parameters, 63
- /etc/ssh/ssh\_host\_key file
  - description, 67
- /etc/ssh/ssh\_known\_hosts file
  - configuring Secure Shell, 64
- /etc/ssh/ssh\_known\_hosts file (*continued*)
  - controlling distribution, 67
  - description, 68
- /etc/ssh/sshd\_config file
  - description, 67
- /etc/ssh/sshrcc file
  - description, 69
- /etc/syslog.conf file
  - PAM, 41
- event
  - description, 326
- event modifier field flags (header token), 382
- events
  - audit
    - See* audit events
  - kernel events
    - auditconfig command options, 361
  - user-level events
    - auditconfig command options, 362
- ex audit flag, 369
- exec\_args token
  - argv policy and, 379
  - format, 379
- exec\_attr database
  - description, 285, 291
  - RBAC relationships, 287
- exec audit class, 369
- exec\_env token
  - format, 379
- execute permissions
  - symbolic mode, 220
- execution attributes
  - description, 291
- execution log (ASET), 299
- exit subcommand
  - sftp command, 58
- exit token
  - format, 380

**F**

- F option
  - allocate command, 392
  - deallocate command, 393
  - st\_clean script, 397
- fa audit flag, 369
- failed login attempts, 238

- failure
  - audit flag prefix, 369
  - turning off audit flags for, 370
- FallBackToRsh keyword
  - ssh\_config file, 65
- fc audit flag, 369
- fd audit flag, 369
- fd\_clean script
  - description, 397
- file\_attr\_acc audit class, 369
- file\_attr\_mod audit class, 369
- file\_close audit class, 369
- file\_creation audit class, 369
- file\_deletion audit class, 369
- file\_read audit class, 369
- file token
  - format, 380
- file vnode token, 378
- file\_write audit class, 369
- files
  - copying with Secure Shell, 57
  - device allocation lock, 353
  - kdc.conf, 187
  - table of SEAM, 181
  - transferring with Secure Shell, 58
- files and file systems
  - ACL entries
    - adding or modifying, 232
    - checking, 231
    - deleting, 202, 233
    - displaying, 202, 233
    - setting, 229, 231
    - valid entries, 227
  - ASET checks, 297
  - ownership
    - changing, 201
    - setgid permission and, 214
    - setuid permission and, 213
  - permissions
    - absolute mode, 219, 221
    - changing, 201, 219, 224
    - defaults, 214
    - description, 212
    - setgid, 214
    - setuid, 213
    - sticky bit, 214
    - symbolic mode, 219, 223
    - umask setting, 214
- files and file systems (*continued*)
  - security, 211, 227
    - access restriction, 198
    - ACL, 201, 227
    - changing ownership, 217
    - changing permissions, 219, 224
    - directory permissions, 212
    - displaying file information, 201, 215
    - encryption, 201
    - file permissions, 212
    - file types, 215
    - special file permissions, 214, 219, 225
    - umask default, 214
    - user classes, 211
  - sharing files, 209
- find command
  - finding files with setuid permissions, 224
- firewall.rpt file, 299
  - description, 302
- firewall systems
  - ASET setup, 207, 299
  - description, 200, 206
  - outside connections with Secure Shell
    - from command line, 60
    - from configuration file, 59
  - packet smashing, 207
  - trusted host, 207
- flags, 371
  - audit
    - See* audit flags
    - audit\_control file line, 364
    - audit\_user file, 366
    - auditconfig command options, 362
    - definitions, 368
    - machine-wide, 364, 368
    - overview, 368
    - prefixes, 370
    - process preselection mask, 371
    - syntax, 369
- flags: line in audit\_control file
  - description, 364
  - prefixes in, 370
  - process preselection mask, 371
- fm audit flag, 369
- forced cleanup, 397
- forwardable tickets
  - definition, 186
  - description, 73

- forwardable tickets (*continued*)
  - example, 174
- forwarding
  - specifying in `ssh_config`, 64
- ForwardX11 keyword
  - Secure Shell port forwarding, 65
- FQDN (Fully Qualified Domain Name)
  - in SEAM, 85
- fr audit flag, 369
- ftp command
  - authentication, 208
- ftp service name
  - PAM, 46
- fw audit flag, 369

## G

- GatewayPorts keyword
  - `ssh_config` file, 64
  - `sshd_config` file, 66
- Generic Security Service API, *See* GSS-API
- get subcommand
  - sftp command, 58
- getclass option
  - auditconfig command, 362
- getcond
  - auditconfig command, 362
- getfacl command
  - description, 202
  - displaying ACL entries, 233
  - examples, 233
  - verifying ACLs set on files, 229
- getpinfo option
  - auditconfig command, 362
- getpolicy option
  - auditconfig command, 362
- getting
  - access to a specific service, 192
  - credential for a server, 191
  - credential for a TGS, 190
- gkadmin command
  - See also* SEAM Administration Tool
  - description, 183
  - `.gkadmin` file, 137
  - description, 181
- GlobalKnownHostsFile keyword
  - `ssh_config` file, 64

- group ACL entries
  - default entries for directories, 228
  - description, 227
  - setting, 229, 231
- group audit policy
  - description, 332
  - groups token and, 332
- group identifier numbers (GIDs)
  - special logins and, 203
- group policy
  - groups token, 381
- group policy
  - groups token and, 381
- group policy
  - newgroups token, 385
- group policy
  - newgroups token and, 385
- groups
  - changing file ownership, 218
- groups token, 381
  - format, 381
- GSS-API
  - SEAM and, 72, 79
- gsscred table
  - using, 193

## H

- hard-disk-space requirements
  - BSM and, 334
- hard string with `audit_warn` script, 367
- header token
  - description, 382
  - event-modifier field flags, 382
  - format, 382
  - order in audit record, 382
- help
  - context-sensitive, 138
  - Help Contents, 138
  - SEAM Administration Tool, 137
  - URL for online, 87
- Help button
  - SEAM Administration Tool, 138
- hierarchical realms
  - configuring, 98
  - in SEAM, 76, 84
- high ASET security level, 296

- host
  - authentication in Secure Shell, 50
  - disabling service on, 171
- Host keyword
  - ssh\_config file, 63
- host names
  - mapping onto realms, 85
- host principal
  - and DNS, 86
  - creating, 94
- HostDSAKey keyword
  - sshd\_config file, 65
- HostKey keyword
  - sshd\_config file, 65
- Hostname keyword
  - ssh\_config file, 65
- hosts
  - trusted host, 207
- hosts.equiv file
  - description, 68

## I

- I option
  - deallocate command, 393
  - st\_clean script, 397
- identity file (Secure Shell)
  - naming convention, 50
- IdentityFile keyword
  - ssh\_config file, 64
- IDs
  - audit, 371
    - overview, 323
  - audit session, 371
  - auditconfig command options, 362
  - mapping UNIX to Kerberos principals, 193
  - terminal, 371
- in\_addr token
  - format, 382
- init service name
  - PAM, 46
- initial ticket
  - definition, 186
- instance
  - in principals names, 76
- integrity
  - SEAM and, 71

- integrity (*continued*)
  - security service, 78
- interactively running ASET, 313
- Internet firewall setup, 200
- Internet-related tokens
  - in\_addr token, 382
  - ip token, 383
  - ipport token, 385
  - socket token, 389
- invalid ticket
  - definition, 186
- io audit flag, 369
- ioctl() system calls, 397
- ioctl audit class, 369
- ioctl system calls, 369
- IP address
  - Secure Shell checking, 64
- ip audit flag, 369
- ip token
  - format, 383
- ipc audit class, 369
- ipc\_perm token
  - format, 384
- ipc token, 384
  - format, 384
- ipc type field values (ipc token), 384
- ipport token
  - format, 385
- item size field
  - arbitrary token, 377

## K

- .k5.REALM file
  - description, 182
- .k5login file
  - description, 181
- kadm5.acl file, 93, 113, 139, 146, 148, 152
  - description, 182
  - format of entries, 152
- kadm5.keytab file, 167
  - description, 182
- kadmin command, 94, 136
  - description, 184
- ktadd command, 168
- ktremove command, 169
- removing principals from keytab with, 169

- kadmin.local command, 93, 141
  - adding administration principals, 93
  - description, 184
- kadmin.log file
  - description, 182
- kadmind daemon
  - master KDC and, 185
  - SEAM and, 184
- kadmind principal, 167
- kdb5\_util command, 92, 97
  - description, 184
- KDC
  - adding entries to propagation file, 94
  - adding slave names to cron job, 97
  - backing up and propagating, 114
  - configuring master, 91
  - configuring server, 90
  - configuring slave, 95
  - copying administration files from slave to master, 97
  - creating database, 92
  - creating host principal, 94
  - creating root principal, 94, 96
  - database propagation, 87
  - master
    - definition, 185
  - planning, 86
  - ports, 86
  - propagating database with
    - kprop\_util, 97
  - restricting access to servers, 120
  - slave, 86
    - definition, 185
  - slave or master, 77, 90
  - starting daemon, 98
  - swapping master and slave, 111
  - synchronizing clocks, 95, 98
- kdc.conf file
  - description, 182
  - ticket lifetime and, 187
- kdc file
  - description, 181
- kdc.log file
  - description, 182
- kdc.master file
  - description, 181
- kdestroy command
  - example, 176
- KeepAlive keyword
  - ssh\_config file, 65
  - sshd\_config file, 66
- KERB authentication
  - dfstab file option, 104
- Kerberos
  - and Kerberos V5, 72
  - and SEAM, 71
  - dfstab file option, 104
  - terminology, 184
- Kerberos (KERB) authentication, 104
- kernel events
  - auditconfig command options, 361
  - BSM and, 327
- key
  - creating for an NIS user, 34
  - creating for Secure Shell, 52
  - description, 185
  - private, 185
  - service, 185
  - service key, 166
  - session, 185, 189
- Key Distribution Center, *See* KDC
- keylogin command, 33
  - running, 29
- KeyRegenerationInterval keyword
  - sshd\_config file, 65
- keyserv daemon
  - starting, 32
  - verifying, 32
- keytab file
  - adding master KDC's host principal to, 94
  - adding service principal to, 167
  - administering, 166
  - administering with ktutil command, 167
  - creating, 93
  - disabling a host's service with
    - delete\_entry command, 172
  - read into keytab buffer with read\_kt command, 170
  - read into keytab with read\_kt command, 172
  - removing principals with ktremove command, 169
  - removing service principal from, 169
  - viewing contents with ktutil command, 169

keytab file (*continued*)  
 viewing keylist buffer with `list`  
 command, 170, 172

kinds of tickets, 186

kinit command  
 example, 174  
 -F option, 174  
 ticket lifetime, 187

klist command  
 example, 174  
 -f option, 174

known\_hosts file  
 configuring Secure Shell, 64  
 controlling distribution, 67  
 description, 68  
 role in authentication, 62

Korn shell  
 ASET working directory specification, 309  
 privileged version, 249

kpasswd command  
 and `passwd` command, 178  
 error message, 178  
 example, 179

kprop command  
 description, 183

kprop\_script script, 97

kpropd.acl file, 94  
 description, 182

kpropd daemon  
 SEAM and, 184

krb5.conf file  
 description, 182  
 domain\_realm section, 85  
 editing, 91  
 ports definition, 86

krb5.keytab file  
 description, 182

krb5 module  
 description, 43

krb5cc\_uid file  
 description, 182

krb5kdc daemon, 98  
 master KDC and, 185  
 SEAM and, 184

ksh command, 240  
 dial-up passwords, 239  
 privileged version, 249

ktadd command, 167

ksh command (*continued*)  
 syntax, 168

ktremove command, 169

ktutil command, 167  
 delete\_entry command, 172  
 list command, 170, 172  
 read\_kt command, 170, 172  
 viewing list of principals, 169

## L

-l option  
 praudit command, 361

-L option  
 ssh command, 56, 64

lcd subcommand  
 sftp command, 58

ldap module  
 description, 43

legacy application  
 securing, 276

lifetime of ticket  
 in SEAM, 187

list command, 170, 172

list\_devices command, 393  
 authorizations required, 294

list privileges in SEAM Administration  
 Tool, 165

ListenAddress keyword  
 sshd\_config file, 66

lo audit flag, 369

LocalForward keyword  
 ssh\_config file, 64

lock files  
 how the allocate mechanism works, 398  
 setting up, 353

log files  
 ASET execution log, 299  
 monitoring su command, 205, 243

logging in  
 displaying user's login status, 235  
 root login  
 access restrictions, 241  
 account, 203, 241  
 restricting to console, 241  
 tracking, 205



- logging in (*continued*)
  - security
    - access restrictions, 202
    - saving failed attempts, 238
    - system access control, 198
    - system device access control, 198
    - tracking root login, 205
  - system logins, 203
- .login file
  - path variable entry, 199
- login file
  - restricting root access to console, 242
  - restricting root access to devices, 241
- login\_logout audit class, 369
- login service name
  - PAM, 46
- logindevperm file
  - description, 198
- LoginGraceTime keyword
  - sshd\_config file, 66
- loginlog file
  - overview, 238
  - saving failed login attempts, 238
- logins command
  - displaying user's login status, 235
  - displaying users with no passwords, 236
  - syntax, 235
- LogLevel keyword
  - ssh\_config file, 65
  - sshd\_config file, 66
- low ASET security level, 296
- ls subcommand
  - sftp command, 58
- lsevent option
  - auditconfig command, 362
- lspolicy option
  - auditconfig command, 362

**M**

- mail
  - using with Secure Shell, 56
- makedbm command
  - description, 292
- managing
  - passwords with SEAM, 176
  - RBAC information, 262
- managing devices, 353
- mapping
  - hostnames onto realms (SEAM), 85
  - UIDs to Kerberos principals, 193
- mappings
  - events to classes (BSM), 327
- mask, process preselection
  - auditconfig command options, 362
  - description, 371
  - machine-wide, 364
- mask ACL entries
  - default entries for directories, 228
  - description, 227
  - setting, 229, 231
- master files
  - ASET, 297, 303
- master KDC
  - configuring, 91
  - definition, 185
  - slave KDCs and, 77, 90
  - swapping with slave KDC, 111
- max\_life value
  - description, 187
- max\_renewable\_life value
  - description, 188
- MaxStartups keyword
  - sshd\_config file, 66
- medium ASET security level, 296
- minfree: line in audit\_control file
  - audit\_warn condition, 367
  - description, 364
- minus (-) audit flag prefix, 370
- minus sign (-)
  - file permissions symbol, 220
- mkdir subcommand
  - sftp command, 58
- modifying
  - policies (SEAM), 160
  - principal (SEAM), 148
  - principal's password, 149
  - roles (RBAC), 268
  - users (RBAC), 274
- module types
  - PAM, 44
- modules
  - PAM, 42
- monitoring
  - audit trail in real time, 334

- monitoring (*continued*)
  - su command use, 205, 243
  - system usage, 199
- mounting
  - NFS file systems, 109
- mounting audit directories, 372
- mt command
  - BSM device cleanup and, 396

## N

- n
  - audit command, 358
- na audit flag, 369
- naflags: line in audit\_control file, 364
- name service scope
  - description, 251
- names
  - audit classes, 368
  - audit files
    - closed files, 373
    - form, 372
    - still-active files, 373
    - time stamps, 373
    - use, 373
  - audit flags, 368
  - audit IDs, 323
  - device names
    - device\_allocate file, 395
    - device\_maps file, 394
  - IDs
    - audit, 371
    - audit session, 371
    - auditconfig command options, 362
    - terminal, 371
  - kernel events, 327
  - user-level events, 327
- naming convention
  - Secure Shell identity file, 50
- ncsd command
  - description, 292
- network audit class, 369
- network security
  - authentication, 208
  - authorization, 208
  - firewall systems
    - description, 200, 206

- network security, firewall systems (*continued*)
  - need for, 200
  - packet smashing, 207
  - trusted host, 207
- issues, 198
- overview, 206
- restricting root access, 209
- Network Time Protocol, *See* NTP
- never-audit flags, 366
  - description, 366
- newgroups token
  - format, 385
  - group policy, 385
- newkey command
  - creating keys for an NIS user, 34
  - generating keys, 29
- NFS
  - mounting systems, 109
- NFS servers
  - ASET and, 307
  - configuring for SEAM, 101
- NFS system, 28
- NIS+
  - ASET checks, 306
  - authentication, 208
  - authorization, 208
  - cred database, 33
  - publickey database, 29
- nisaddcred command, 33
  - generating keys, 29
- no audit flag, 369
- no\_class audit class, 369
- nobody user, 209
- nologin file
  - description, 68
- non\_attrib audit class, 369
- non-hierarchical realms
  - in SEAM, 76
- nonattributable flags in audit\_control file, 364
- not\_terminated files
  - cleaning, 351
- nt audit flag, 369
- NTP, 95, 98
  - SEAM and, 87
- null audit class, 369
- NumberOfPasswordPrompts keyword
  - ssh\_config file, 65

## O

- O option
  - auditreduce command, 350
- object-reuse requirement, 397
  - device-clean scripts
    - audio devices, 397
    - CD-ROM drives, 397
    - described, 396
    - diskette drives, 397
    - tape drives, 395
    - writing new scripts, 397
  - in BSM, 396
- obtaining
  - access to a specific service, 192
  - credential for a server, 191
  - credential for a TGS, 190
  - forwardable tickets, 174
  - tickets with kinit, 174
- online help
  - context-sensitive, 138
  - Help Contents, 138
  - SEAM Administration Tool, 137
  - URL for, 87
- opaque token
  - format, 386
- Operator rights profile
  - description, 280, 282
- Operator role
  - description, 246
- option control flag
  - PAM, 47
- ot audit flag, 369
- other ACL entries
  - default entries for directories, 228
  - description, 227
  - setting, 229, 231
- other audit class, 369
- overflow prevention
  - audit trail, 352
- ovsec\_admin.xxxx file
  - description, 182
- ownership of files
  - ACLs and, 201, 227
  - changing, 201, 217
  - changing group ownership, 218

## P

- packet transfers
  - firewall security, 200
  - packet smashing, 207
- PAM
  - add a module, 40
  - configuration file, 44, 47, 182
  - control flags, 46
  - /etc/syslog.conf file, 41
  - module types, 44
  - modules, 42
  - overview, 36
  - password mapping, 38
  - planning, 39
  - SEAM and, 79, 183
  - service names, 45
  - stacking, 38
  - try\_first\_pass, 179
- pam\_\*.so.1 files
  - description, 42
- pam.conf file
  - description, 182
  - SEAM and, 182
- pam\_roles command
  - description, 293
- panels, table of SEAM Administration Tool, 162
- passphrase
  - example, 53
- passwd command
  - and kpasswd command, 178
  - try\_first\_pass, 179
- passwd file
  - ASET checks, 298
  - /etc/d\_passwd file and, 239
- passwd service name
  - PAM, 46
- password mapping
  - in PAM, 38
- PasswordAuthentication keyword
  - sshd\_config file, 65
- passwords
  - and policies, 178
  - capturing encrypted passwords, 241
  - changing with kpasswd command, 178
  - changing with passwd command, 178
  - dial-up passwords
    - disabling dial-up logins temporarily, 241

- passwords, dial-up passwords (*continued*)
  - /etc/d\_passwd file, 239
  - /etc/dialups file, 239
- displaying users with no passwords, 236
- eeprom security, 203
- eliminating in Secure Shell use, 54
- login security, 198, 202
- management, 176
- modifying a principal's password, 149
- secret-key decryption, 29
- Secure Shell, 50
- suggestions on choosing, 177
- system logins, 203
- UNIX and Kerberos, 176
- path audit policy
  - description, 332
- path token, 386
- path variable
  - setting, 199
- pc audit flag, 369
- PERIODIC\_SCHEDULE variable (ASET)
  - scheduling ASET, 306, 310, 314
- permissions
  - ACLs and, 201, 227
  - ASET handling of, 296
  - changing file permissions
    - absolute mode, 219, 221
    - chmod command, 201
    - symbolic mode, 219, 223
  - defaults, 214
  - directory permissions, 212
  - file permissions
    - absolute mode, 219, 221
    - changing, 219, 224
    - description, 212
    - special permissions, 214, 219, 225
    - symbolic mode, 219, 223
  - setgid permissions
    - absolute mode, 219, 223
    - description, 214
    - symbolic mode, 220
  - setuid permissions
    - absolute mode, 219, 223
    - description, 213
    - finding files with permissions set, 224
    - security risks, 213
    - symbolic mode, 220
  - special file permissions, 214, 219, 225
- permissions (*continued*)
  - sticky bit, 214
  - tune files (ASET), 303, 306
  - umask settings, 214
  - user classes and, 212
- PermitEmptyPasswords keyword
  - sshd\_config file, 66
- PermitRootLogin keyword
  - sshd\_config file, 66
- pfcsch command
  - description, 249
- pfexec command
  - description, 293
- pfksh command
  - description, 249
- pfsh command
  - description, 249
- physical security, 198
- planning
  - PAM, 39
  - RBAC, 254
  - SEAM
    - client and service principal names, 85
    - clock synchronization, 87
    - configuration decisions, 83
    - database propagation, 87
    - number of realms, 84
    - ports, 86
    - realm hierarchy, 84
    - realm names, 84
    - realms, 84
    - slave KDCs, 86
- pluggable authentication module, *See* PAM
- plus (+) audit flag prefix, 370
- plus sign (+)
  - file permissions symbol, 220
- policies
  - administering, 135, 153
  - and passwords, 178
  - auditconfig options, 362
  - creating (SEAM), 146
  - creating new (SEAM), 158
  - deleting, 161
  - modifying, 160
  - SEAM Administration Tool panels for, 162
  - task map for administering, 153
  - viewing attributes, 156
  - viewing list of, 154

- policy.conf database
  - Basic Solaris User rights profile, 283
  - description, 292
  - RBAC relationships, 287
- port
  - for KDC and admin services, 86
  - KDC administration daemon, 86
- port forwarding
  - configuring ssh\_config, 64
  - Secure Shell, 56
- Port keyword
  - sshd\_config file, 66
- postdatable ticket
  - definition, 187
- postdated ticket
  - description, 73
- postsigterm string
  - audit\_warn script, 368
- pound sign (#)
  - device\_allocate file, 395
  - device\_maps file, 394
- ppp service name
  - PAM, 46
- praudit command
  - converting audit records to readable format, 328, 359
  - output formats, 360
  - piping auditreduce output to, 349
  - using, 360
- prefixes in audit flags, 370
- preselection mask
  - auditconfig command options, 362
  - description, 371
  - machine-wide, 364
- preselection mask (BSM)
  - reducing storage costs, 334
- primary
  - in principals names, 76
- Primary Administrator
  - rights profile, 280, 282
  - role, 246
- primary audit directory, 364
- principal
  - adding administration, 93
  - adding service principal to keytab, 167
  - administering, 135, 140
  - automating creation of, 141
  - creating host, 94
- principal (*continued*)
  - creating root, 94, 96
  - deleting, 149
  - duplicating, 148
  - in SEAM, 76
  - modifying, 148
  - principal name, 76
  - removing from keytab file, 169
  - removing service principal from keytab, 169
  - root, 94
  - SEAM Administration Tool panels for, 162
  - service principal, 76
  - setting up defaults, 150
  - task map for administering, 141
  - user ID comparison, 103
  - user principal, 76
  - viewing attributes, 144
  - viewing list of, 142
  - viewing sublist of principals, 143
- principal.db file
  - description, 182
- principal.kadm5 file
  - description, 182
- principal.kadm5.lock file
  - description, 182
- principal.ok file
  - description, 182
- principals
  - creating, 146
- print format field
  - arbitrary token, 377
- Printer Management rights profile
  - description, 280, 283
- printing
  - audit log, 349
- privacy
  - SEAM and, 71
  - security service, 78
- private key, 28
  - definition in SEAM, 185
  - description, 50
  - naming convention, 50
- privilege, 165
  - effects on SEAM Administration Tool, 166
- privileged application
  - authorization checking, 249
  - description, 246

- privileged application (*continued*)
  - ID checking, 249
- process audit characteristics
  - audit ID, 371
  - audit session ID, 371
  - process preselection mask, 371
  - terminal ID, 371
- process audit class, 369
- process preselection mask
  - auditconfig command options, 362
  - description, 371
- process token
  - format, 387
- processing time costs
  - BSM and, 333
- prof\_attr database
  - description, 285, 290
  - RBAC relationships, 287
- profile, *See* rights profile
- .profile file
  - path variable entry, 199
- profile shell
  - description, 249
- profiles command
  - description, 293
- program
  - testing for authorizations, 276
- projects module
  - description, 43
- propagation
  - KDC database, 87
  - Kerberos database, 114
- propagation file
  - adding entries to, 94
- Protocol keyword
  - sshd\_config file, 66
- proxiable ticket
  - definition, 187
- proxy ticket
  - definition, 187
- ProxyCommand keyword
  - ssh\_config file, 65
- pseudo-tty
  - use in Secure Shell, 62
- public directories, 214
- public key
  - description, 50
  - DH authentication and, 28

- public key (*continued*)
  - known hosts file, 64
  - naming convention, 50
  - Secure Shell, 50
- public-key cryptography
  - AUTH\_DH client-server session, 29, 32
  - changing public and secret keys, 29
  - common key
    - calculation, 31
  - database of public keys, 29
  - generating keys
    - conversation key, 30
    - public and secret keys, 29
  - secret key
    - changing, 29
    - database, 29
    - decrypting, 29
    - generating, 29
- publickey map
  - DH authentication and, 28
- put subcommand
  - sftp command, 58

## Q

- question mark (?) wildcard character
  - in ASET tune files, 312
- quit subcommand
  - sftp command, 58

## R

- R option
  - ssh command, 56, 64
- r praudit output format, 361
- raw praudit output format, 361
- RBAC
  - administration commands, 292
  - authorization database, 288
  - basic concept, 245
  - database relationships, 287
  - elements, 246
  - name services, 287
  - rights profile database, 290
  - tasks, 254
    - adding custom roles, 267

- RBAC, tasks (*continued*)
  - adding first role, 259
  - adding first user, 257
  - adding rights profile example, 273
  - adding roles, 265
  - adding roles from command line, 267
  - changing rights profiles from command line, 274
  - changing roles from command line, 270
  - changing user properties from command line, 275
  - checking scripts or programs for
    - authorizations, 276
  - configuration, 254
  - editing rights profiles, 270
  - information management task map, 262
  - modifying roles, 268
  - modifying users, 274
  - planning, 254
  - running the user tools, 256
  - securing legacy applications, 276
  - securing scripts, 276
  - setting IDs on commands, 272
  - using privileged applications, 263
- rc file
  - description, 69
- rcp command
  - authentication, 208
- read into keytab buffer with `read_kt`
  - command, 170
- read into keytab with `read_kt` command, 172
- `read_kt` command, 170, 172
- read permissions
  - symbolic mode, 220
- readable audit record format
  - converting audit records to, 328, 359
- reallocating devices, 392
- realms
  - and servers, 77
  - configuration decisions, 84
  - configuring cross-realm authentication, 98
  - contents of, 77
  - direct, 99
  - hierarchical, 98
  - hierarchical or non-hierarchical, 76
  - hierarchy, 84
  - in principal names, 76
  - in principals names, 76
- realms (*continued*)
  - mapping hostnames onto, 85
  - names, 84
  - number of, 84
- reducing
  - audit files, 350
  - storage-space requirements for audit files, 334
- reducing audit files
  - `auditreduce` command, 358, 360
- remote logins
  - authentication, 208
  - authorization, 208
  - security and, 31
- remote systems
  - logging in
    - authentication, 208
    - authorization, 208
- removing
  - principals with `ktremove` command, 169
  - service principal from keytab file, 169
- renewable ticket
  - definition, 187
- replayed transactions, 31
- reports
  - ASET, 302, 308
- reports directory (ASET), 302
- required control flag
  - PAM, 46
- requisite control flag
  - PAM, 47
- restoring
  - ASET, 307
- restricted shell (`rsh`), 204
- restricting access for KDC servers, 120
- return token
  - format, 388
- `rewoffl` option
  - `mt` command
    - BSM device cleanup and, 396
- `rexd` service name
  - PAM, 46
- `.rhosts` file
  - description, 68
  - role in authentication, 62
- `rhosts` module
  - description, 43

- RhostsAuthentication keyword
  - sshd\_config file, 65
- RhostsRSAAuthentication keyword
  - sshd\_config file, 65
- right, *See* rights profile
- rights profile
  - See also* individual profiles
  - changing rights profiles from command line, 274
  - creation example, 273
  - database
    - See* prof\_attr database and exec\_attr database
  - description, 246, 251
  - editing, 270
  - major rights profiles description, 280
- Rights tool
  - description, 270
- rlogin command
  - authentication, 208
- rlogin service name
  - PAM, 46
- role
  - adding custom roles, 267
  - adding first role, 259, 262
  - adding roles, 265
  - adding roles from command line, 267
  - assuming, 250
  - assumption example, 263
  - changing roles from command line, 270
  - description, 246, 250
  - making root a role, 261
  - modifying roles, 268
  - properties
    - summarized, 270
  - recommended role rights profiles, 279
  - recommended roles, 246
  - use in RBAC, 245
- role-based access control, *See* RBAC
- Role Properties dialog box
  - description, 269
- roleadd command
  - description, 293
- roledel command
  - description, 293
- rolemod command
  - description, 293
- roles command
  - description, 293
- roles module
  - description, 43
- root
  - adding principal to host's keytab, 167
  - authentication for NFS, 109
  - eliminating root in RBAC, 250
- root access
  - displaying attempts on console, 243
  - monitoring su command use, 205, 243
  - restricting, 209, 241
- root login
  - account
    - access restrictions, 241
    - description, 203
    - restricting to console, 242
    - tracking, 205
- root principal
  - creating, 94, 96
- root role
  - creating, 261
- RPCSEC\_GSS API
  - SEAM and, 80
- RSAAuthentication keyword
  - sshd\_config file, 65
- rsh command (restricted shell), 204
- rsh service name
  - PAM, 46
- running the User tool
  - task description, 256

**S**

- s
  - audit command, 358
  - praudit command, 361
- S option of st\_clean script, 398
- sac service name
  - PAM, 46
- sample module
  - description, 43
- saving
  - failed login attempts, 238
- scheduling ASET execution
  - (PERIODIC\_SCHEDULE), 296, 306, 310, 314



- scope
  - description, 251
- scp command
  - authentication steps, 62
  - description, 69
  - using, 57
- script
  - securing, 276
  - testing for authorizations, 276
- SCSI devices
  - st\_clean script, 395
- SEAM
  - administering, 135
  - Administration Tool, 136
  - and Kerberos V5, 71
  - commands, 183
  - components of, 79
  - configuration decisions, 83
  - configuring KDC servers, 90
  - daemons, 184
  - files, 181
  - gaining access to server, 190
  - online help, 87
  - overview, 72
  - overview of authentication, 189
  - password management, 176
  - planning for, 83
  - reference, 181
  - terminology, 184
  - using, 173
- SEAM Administration Tool, 136
  - and limited administration privileges, 165
  - and list privileges, 165
  - and X Window system, 137
  - command-line equivalents, 137
  - context-sensitive help, 138
  - creating a new principal, 146
  - creating new policy, 146, 158
  - default values, 139
  - deleting a principal, 149
  - deleting policies, 161
  - displaying sublist of principals, 143
  - duplicating a principal, 148
  - files modified by, 137
  - Filter Pattern field, 143
  - gkadmin command, 135
  - gkadmin command vs. kadmin, 136
  - .gkadmin file, 137
- SEAM Administration Tool (*continued*)
  - help (print), 137
  - Help button, 138
  - Help Contents, 138
  - how affected by privileges, 166
  - kadmin command vs. gkadmin, 136
  - login window, 139
  - modifying a principal, 148
  - modifying policies, 160
  - online help, 137
  - panel descriptions, 162
  - privileges, 165
  - setting up principal defaults, 150
  - starting, 139
  - table of panels, 162
  - viewing a principal's attributes, 144
  - viewing list of policies, 154
  - viewing list of principals, 142
  - viewing policy attributes, 156
  - vs. kadmin command, 136
- searching
  - files with `setuid` permissions, 224
- secondary audit directory, 364
- secret key
  - changing, 29
  - database, 29
  - decrypting, 29
  - generating, 29
- secure access, 35
- secure NIS+, adding a user, 34
- Secure RPC, 27
  - implementation of, 29
- Secure RPC authentication, 208
- Secure Shell
  - administering, 61
  - authentication, 50
  - authentication steps, 62
  - configuring, 63
  - configuring clients, 63
  - connecting outside firewall
    - from command line, 60
    - from configuration file, 59
  - copying files, 57
  - creating keys, 52
  - description, 49
  - forwarding mail, 56
  - important files, 67
  - local port forwarding, 56

- Secure Shell (*continued*)
  - logging in, 53
  - naming identity files, 50
  - port forwarding, 56
  - protocol versions, 49
  - public key, 50
  - remote port forwarding, 57
  - transferring files, 58
  - typical session, 61
  - user task map, 52
  - using without password, 54
- securing legacy applications
  - description, 276
- securing scripts
  - description, 276
- security
  - auditing and, 325
  - DH authentication
    - AUTH\_DH client-server session, 29, 32
    - KERB authentication, 104
  - security mode
    - setting up environment with multiple, 104
  - security service
    - in SEAM, 78
    - integrity, 78
    - privacy, 78
  - seq audit policy
    - description, 332
    - seq token and, 332
  - seq policy
    - seq token and, 388
  - seq token
    - format, 388
    - seq policy and, 388
- server authentication parameters
  - sshd\_config file, 65
- ServerKeyBits keyword
  - sshd\_config file, 66
- servers
  - and realms, 77
  - AUTH\_DH client-server session, 29, 32
  - configuring for Secure Shell, 65
  - definition in SEAM, 185
  - gaining access with SEAM, 190
  - obtaining credential for, 191
- service
  - definition in SEAM, 185
  - disabling on a host, 171
- service (*continued*)
  - obtaining access for specific service, 192
- service key, 166
  - definition in SEAM, 185
- service names
  - PAM, 45
- service principal
  - adding to keytab file, 167
  - description, 76
  - planning for names, 85
  - removing from keytab file, 169
- session ID, 371
- session key
  - definition in SEAM, 185
  - SEAM authentication and, 189
- setclass option
  - auditconfig command, 362
- setcond option
  - auditconfig command, 362
- setenv command
  - ASET security level specification, 309
  - ASET working directory specification, 309
- setfacl command
  - adding ACL entries, 232
  - deleting ACL entries, 233
  - description, 202
  - examples, 230, 232
  - modifying ACL entries, 232
  - setting ACL entries, 229, 231
  - syntax, 229
- setgid permissions
  - absolute mode, 219, 223
  - description, 214
  - symbolic mode, 220
- setpmask option of auditconfig command
  - auditconfig command, 362
- setpolicy option
  - auditconfig command, 362
- setsmask option
  - auditconfig command, 362
- setting IDs on commands
  - description, 272
  - task description, 272
- setting up principal defaults, 150
- setuid permissions
  - absolute mode, 219, 223
  - description, 213
  - finding files with permissions set, 224

- setuid permissions (*continued*)
  - security risks, 213
  - symbolic mode, 220
- setuid programs, 199
- setumask option
  - auditconfig command, 362
- sftp command
  - authentication steps, 62
  - description, 69
  - using, 58
- sh command, 240
  - dial-up passwords, 239
  - privileged version, 249
- share command
  - restricting root access, 209
- sharing files (network security), 209
- shell
  - privileged versions, 249
- shell commands
  - /etc/d\_passwd file entries, 239
- shell programs
  - ASET security level specification, 309
  - ASET working directory specification, 309
- short praudit output format, 361
- shosts.equiv file
  - description, 68
- .shosts file
  - description, 68
- signal received during auditing shutdown, 368
- single-sign-on system
  - SEAM and, 71
- size
  - reducing audit files, 350
    - auditreduce command, 358
    - auditreduce command, 360
  - reducing storage-space requirements for audit files, 334
- slave\_datatrans file, 114
  - description, 182
- slave KDCs
  - adding names to cron job, 97
  - configuring, 95
  - definition, 185
  - master KDC and, 77
  - or master, 90
  - planning for, 86
  - swapping with master KDC, 111
- smartcard module
  - description, 43
- smattrpop command
  - description, 293
- SMC, *See* Solaris Management Console
- smexec command
  - description, 293
- smmultiuser command
  - description, 293
- smprofile command
  - description, 293
- smrole command
  - description, 293
- smuser command
  - description, 293
- socket token, 389
- soft limit
  - audit\_warn condition, 367
  - minfree: line description, 364
- soft string with audit\_warn script, 367
- Solaris Management Console
  - role assumption, 264
  - running the user tools, 256
- sr\_clean script
  - description, 397
- ssh-add command
  - description, 69
  - example, 54
- ssh-agent command
  - description, 69
  - from command line, 54
  - in scripts, 55
- ssh command
  - authentication steps, 62
  - description, 69
  - L option, 56
  - o option, 60
  - permitting access, 66
  - port forwarding, 64
  - R option, 56
  - using, 53
- ssh\_config file
  - client authentication parameters, 63
  - configuring Secure Shell, 63
  - connection parameters, 65
  - host-specific parameters, 63

- ssh\_config file (*continued*)
  - keywords
    - See* specific keyword
  - known host file parameters, 64
- ssh\_host\_key file
  - description, 67
- ssh\_host\_key.pub file
  - description, 67
- ssh-keygen command
  - description, 69
  - using, 52
- ssh\_known\_hosts file
  - configuring Secure Shell, 64
  - description, 68
- ssh service name
  - PAM, 46
- sshd command
  - configuring for forwarding, 66
  - description, 69
  - session controls, 66
- sshd\_config file
  - description, 67
  - forwarding parameters, 66
  - ports parameters, 66
  - server connection parameters, 66
  - session control parameters, 66
- sshd.pid file
  - description, 68
- sshrd file
  - description, 69
- st\_clean script
  - description, 396
- st\_clean script for tape drives, 395
- stacking
  - in PAM, 38
- standard cleanup, 398
- starting
  - ASET
    - initiating sessions from shell, 296
    - running interactively, 313
  - KDC daemon, 98
- stash file
  - creating, 97
  - definition, 185
- sticky bit permissions
  - absolute mode, 219, 223
  - description, 214
  - symbolic mode, 220
- stopping
  - dial-up logins temporarily, 241
- storage
  - audit records and, 330
- storage costs
  - BSM and, 334
- storage overflow prevention
  - audit trail, 352
- StrictHostKeyChecking keyword
  - ssh\_config file, 64
- StrictModes keyword
  - sshd\_config file, 66
- su command
  - displaying use on console, 243
  - in role assumption, 263
  - monitoring use, 205, 243
- su file
  - monitoring su command, 243
- su service name
  - PAM, 46
- subject token
  - format, 390
- Subsystem keyword
  - sshd\_config file, 66
- success
  - audit flag prefix, 369
  - turning off audit flags for, 370
- sufficient control flag
  - PAM, 47
- sulog file, 205, 243
- superuser
  - eliminating superuser in RBAC, 250
  - model versus RBAC, 245
- suser
  - security policy, 285
- swapping master and slave KDCs, 111
- symbolic links
  - file permissions, 212
  - latest directory (ASET), 302
- symbolic mode
  - changing file permissions, 220, 223
  - description, 219
- synchronizing clocks, 95, 98, 109
- sysconf.rpt file
  - description, 298, 302
- SyslogFacility keyword
  - sshd\_config file, 66

- System Administrator
  - rights profile, 280, 282
  - role, 246
- system calls
  - arg token, 378
  - auditsvc () fails, 391
  - auditsvc () fails, 368
  - close, 369
  - event numbers, 327
  - exec\_args token, 379
  - exec\_env token, 379
  - ioctl, 369, 397
  - return token, 388
- system security
  - dial-up passwords, 239
    - disabling dial-up logins temporarily, 241
    - /etc/d\_passwd file, 239
    - /etc/dialups file, 239
  - displaying
    - user's login status, 235
    - users with no passwords, 236
  - introduction, 235
  - login access restrictions, 202
  - overview, 202
  - passwords, 203
  - restricted shell, 204
  - restricting root login to console, 241
  - root access restrictions, 209, 241
  - saving failed login attempts, 238
  - special logins, 203
  - su command monitoring, 205, 243
- System V IPC
  - ipc audit class, 369
  - ipc\_perm token, 384
  - ipc token, 384

**T**

- tables
  - gsscred, 193
- tail command
  - BSM and, 335
- tape drives
  - device-clean scripts, 396
  - st\_clean script, 395
- task map
  - administering policies, 153
  - task map (*continued*)
    - administering principals, 141
    - Secure Shell, 52
  - TASKS variable (ASET)
    - configuring ASET, 305, 310
  - taskstat command (ASET), 297, 300
  - TCP address, 385
  - TCP/IP
    - specifying in sshd\_config, 66
    - use in Secure Shell, 62
  - telnet service name
    - PAM, 46
  - temporary file cannot be used, 368
  - terminal ID, 371
  - terminating
    - signal received during auditing shutdown, 368
  - terminology
    - authentication-specific, 185
    - Kerberos-specific, 184
    - SEAM, 184
  - text token, 391
  - TGS
    - getting credential for, 190
  - TGT
    - in SEAM, 73
  - ticket file, *See* credential cache
  - ticket-granting service, *See* TGS
  - Ticket-Granting Ticket, *See* TGT
  - tickets
    - creating, 173
    - creating with kinit, 174
    - definition, 72
    - description, 185
    - destroying, 176
    - file
      - See* credential cache
    - forwardable, 73, 174, 186
    - initial, 186
    - invalid, 186
    - klist command, 174
    - lifetime, 187
    - maximum renewable lifetime, 188
    - obtaining, 173
    - or credentials, 73
    - postdatable, 187
    - postdated, 73
    - proxiabile, 187

- tickets (*continued*)
  - proxy, 187
  - renewable, 187
  - types of, 186
  - viewing, 174
  - warning about expiration, 108
- time stamps in audit files, 373
- /tmp/krb5cc\_uid file
  - description, 182
- /tmp/ovsec\_admin.xxxxx file
  - description, 182
- tmpfile string
  - audit\_warn script, 368
- tmpfs file system, 214
- trail audit policy
  - description, 333
  - trailer token and, 333
- trailer token
  - description, 391
  - format, 391
  - order in audit record, 391
  - praudit display, 391
- transferring files
  - using Secure Shell, 58
- transparency
  - definition in SEAM, 73
- Trojan horse, 199
- trusted host, 207
- try\_first\_pass, 179
- ttymon service name
  - PAM, 46
- tune files (ASET)
  - description, 303, 306
  - example files, 311
  - format, 311
  - modifying, 306
  - rules, 311
- tune.rpt file
  - description, 297, 302
- types of tickets, 186

## U

- U option
  - allocate command, 392
  - list\_devices command, 393
- UDP address, 385

- uid\_aliases file
  - description, 303
  - specifying, 306
- UID\_ALIASES variable (ASET)
  - aliases file specification, 306, 311
  - description, 303
- umask setting, 214
- unix\_account module
  - description, 44
- unix\_auth module
  - description, 44
- unix module
  - description, 43
- unix\_session module
  - description, 44
- URL for online help, 87
- UseLogin keyword
  - sshd\_config file, 66
- user
  - adding first user, 257
  - assigning RBAC defaults, 292
  - changing user properties from command line, 275
  - database
    - See user\_attr database
    - modifying properties, 274
  - user accounts
    - ASET check, 298
    - displaying login status, 235
  - User Accounts tool
    - description, 274
  - user ACL entries
    - default entries for directories, 228
    - description, 227
    - setting, 229, 231
  - user\_attr database
    - description, 285, 287
    - RBAC relationships, 287
  - user audit fields, 366
  - user classes of files, 211
  - user ID
    - audit ID and, 323
    - in NFS services, 103
  - user ID (audit ID), 371
  - User keyword
    - ssh\_config file, 65
  - user-level events
    - auditconfig command options, 362

- user-level events (*continued*)
  - BSM and, 327
- user principal
  - description, 76
- useradd command
  - description, 293
- userdel command
  - description, 293
- UserKnownHostsFile keyword
  - ssh\_config file, 64
- usermod command
  - description, 293
- UseRsh
  - ssh\_config file, 64
- using privileged applications
  - task description, 263
- /usr/aset/asetenv file, 304
  - modifying, 304
  - running ASET periodically, 314
- /usr/aset directory, 296
- /usr/aset/masters/tune files, 303
  - example files, 311
  - format, 311
  - modifying, 306
  - rules, 311
- /usr/aset/masters/uid\_aliases file, 303
- /usr/aset/reports directory
  - structure, 300, 302
- /usr/aset/reports/latest directory, 302
- /usr/lib/krb5/kadmind daemon
  - SEAM and, 184
- /usr/lib/krb5/kprop command
  - description, 183
- /usr/lib/krb5/kpropd daemon
  - SEAM and, 184
- /usr/lib/krb5/krb5kdc daemon
  - SEAM and, 184
- /usr/sbin/gkadmin command
  - description, 183
- /usr/sbin/kadmin command
  - description, 184
- /usr/sbin/kadmin.local command
  - description, 184
- /usr/sbin/kdb5\_util command
  - description, 184
- usrgrp.rpt file
  - description, 298, 302
- usrgrp.rpt file (*continued*)
  - example, 302
- uucico command, 239
- uucp service name
  - PAM, 46

**V**

- v1 protocol
  - Secure Shell, 49
- v2 protocol
  - Secure Shell, 49
- /var/adm/loginlog file
  - overview, 238
  - saving failed login attempts, 238
- /var/adm/sulog file, 205, 243
- /var/krb5/.k5.REALM file
  - description, 182
- /var/krb5/kadmin.log file
  - description, 182
- /var/krb5/kdc.log file
  - description, 182
- /var/krb5/principal.db file
  - description, 182
- /var/krb5/principal.kadm5 file
  - description, 182
- /var/krb5/principal.kadm5.lock file
  - description, 182
- /var/krb5/principal.ok file
  - description, 182
- /var/krb5/slave\_datatrans file
  - description, 182
- /var/run/sshd.pid file
  - description, 68
- variables
  - ASET environment variables
    - ASETDIR, 309
    - ASETSECLEVEL, 309
    - CKLISTPATH\_level, 303, 305, 311
    - PERIODIC\_SCHEDULE, 306, 310, 314
    - summary table, 308
    - TASKS, 305, 310
    - UID\_ALIASES, 303, 306, 311
    - YPCHECK, 306, 311
- verifiers
  - description, 30
  - returned to client, 31

verifiers (*continued*)  
  window, 30  
viewing  
  keylist buffer with `list` command, 170, 172  
  list of policies, 154  
  list of principals, 142  
  policy attributes, 156  
  principal's attributes, 144  
  tickets, 174  
viruses  
  Trojan horse, 199  
vnode token  
  format, 378

## W

warn.conf file  
  description, 182  
warning about ticket expiration, 108  
wildcard characters  
  in ASET tune files, 312  
window verifier, 30  
write permissions  
  symbolic mode, 220  
writing new device-clean scripts, 397

## X

X Window system, and SEAM Administration  
  Tool, 137  
X11  
  use in Secure Shell, 62  
X11 forwarding  
  configuring `ssh_config`, 65  
xauth command  
  X11 forwarding, 65  
XAuthLocation keyword  
  Secure Shell port forwarding, 65  
  `sshd_config` file, 66  
Xylogics tape drive clean script, 395

## Y

YPCHECK variable (ASET)  
  specifying system configuration file  
  tables, 306, 311