Mac OS X Server

Security Configuration
For Version 10.4 or Later
Second Edition

Apple Inc.
© 2007 Apple Inc. All rights reserved.

The owner or authorized user of a valid copy of Mac OS X Server software may reproduce this publication for the purpose of learning to use such software. No part of this publication may be reproduced or transmitted for commercial purposes, such as selling copies of this publication or for providing paid-for support services.

Every effort has been made to ensure that the information in this manual is accurate. Apple Inc., is not responsible for printing or clerical errors.

Apple
1 Infinite Loop
Cupertino CA 95014-2084
408-996-1010
www.apple.com

The Apple logo is a trademark of Apple Inc., registered in the U.S. and other countries. Use of the "keyboard" Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

Apple, the Apple logo, Airport, AppleShare, AppleTalk, FireWire, Keychain, Mac, Macintosh, Mac OS, QuickTime, WebObjects, Xgrid, and Xserve are trademarks of Apple Inc., registered in the U.S. and other countries. Apple Remote Desktop and Finder are trademarks of Apple Inc.

Adobe and PostScript are trademarks of Adobe Systems Incorporated.

The Bluetooth® word mark and logos are owned by the Bluetooth SIG, Inc. and any use of such marks by Apple Inc. is under license.

Intel and Intel Core are trademarks of Intel Corp. in the U.S. and other countries.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

Other company and product names mentioned herein are trademarks of their respective companies. Mention of third-party products is for informational purposes only and constitutes neither an endorsement nor a recommendation. Apple assumes no responsibility with regard to the performance or use of these products.

019-0923/02-15-07

Contents

Preface

15 **About This Guide**
15 Target Audience
15 What's New in Version 10.4
16 What's in This Guide
18 Using This Guide
18 Using Onscreen Help
19 The Mac OS X Server Suite
20 Getting Documentation Updates
21 Getting Additional Information
22 Acknowledgments

Chapter 1

23 **Introducing Mac OS X Server Security Architecture**
24 Security Architecture Overview
24 UNIX Infrastructure
24 Access Permissions
24 Security Framework
25 Layered Security Defense
26 Built-In Security Services
26 Keychain Services
26 Secure Transport Services
26 Certificate, Key, and Trust Services
26 Authorization Services
27 Smart Card Services
27 Directory Services
27 Open Directory Authentication Architecture
28 Policy Management
28 Authorization versus Authentication
29 Network Deployment Considerations
29 Network Isolation
29 Functional Separation
30 Using Accounts Securely

Chapter 2
31 **Installing Mac OS X Server**
31 System Installation Overview
32 Disabling the Open Firmware Password
33 Installing Locally from CD or DVD
34 Installing Remotely from Disks or Images
35 Installing Remotely from the Installation Discs
36 Installing Remotely from an Image
37 Installing from the Command Line
38 Using the installer Command Tool for Installation
38 Using the asr Command Tool for Installation
39 Initializing Server Setup
39 Using Server Assistant
39 Setting Up a Secure Local Server
41 Setting Up a Secure Remote Server
42 Updating System Software
43 Updating from an Internal Software Update Server
43 Updating from Internet-Based Software Update Servers
44 Updating Manually from Installer Packages
45 Verifying the Integrity of Software
45 Repairing Disk Permission
45 Kinds of Permissions
46 POSIX Permissions Overview
46 ACL Permissions Overview
46 Using Disk Utility to Repair Disk Permissions

Chapter 3
49 **Protecting Hardware and Securing Global System Settings**
49 Protecting Hardware
50 Disabling Hardware
51 Removing Mac OS 9
52 Using the Command Line to Remove Mac OS 9
52 Running Mac OS 9 from a CD or DVD
53 Running Mac OS 9 from a Disc Image
54 Securing System Startup
55 Using the Open Firmware Password Application
56 Configuring Open Firmware Settings
58 Using Command-Line Tools to Secure Startup
58 Requiring a Password for Single-User Mode
59 Configuring Access Warnings
59 Enabling Access Warnings for the Login Window
60 Enabling Access Warnings for the Command Line
61 Securing Fast User Switching
61 Displaying a Login Warning Banner
61 Setting a Local Login Warning Banner

62 Setting a Login Warning Banner for Remote Services

Chapter 4
63 **Securing Local Server Accounts**
63 Types of User Accounts
64 General Guidelines for Securing Accounts
64 Defining User IDs
67 Securing Local Nonadministrator Accounts
69 Securing Local Server Administrator Accounts
70 Securing a Local Directory Domain Administrator Account
70 Securing the Local System Administrator Account
71 Restricting sudo Usage
72 Using Strong Authentication
73 Using Password Assistant
74 Using Smart Cards
74 Using Tokens
74 Using Biometrics
75 Storing Credentials
76 Using the Default User Keychain
77 Securing Keychain Items
77 Creating Additional Keychains
79 Using Portable and Network-Based Keychains

Chapter 5
81 **Securing System Preferences**
83 Securing .Mac Preferences
85 Securing Accounts Preferences
88 Securing Appearance Preferences
89 Securing Bluetooth Preferences
90 Securing CDs & DVDs Preferences
90 Securing Classic Preferences
93 Securing Dashboard and Exposé Preferences
94 Securing Date & Time Preferences
95 Securing Desktop & Screen Saver Preferences
97 Securing Displays Preferences
97 Securing Dock Preferences
98 Securing Energy Saver Preferences
99 Securing International Preferences
99 Securing Keyboard & Mouse Preferences
100 Securing Network Preferences
102 Securing Print & Fax Preferences
103 Securing QuickTime Preferences
104 Securing Security Preferences
105 Securing Sharing Preferences
106 Securing Software Update Preferences

107 Securing Sound Preferences
108 Securing Speech Preferences
109 Securing Spotlight Preferences
111 Securing Startup Disk Preferences
112 Securing Universal Access Preferences

Chapter 6 113 **Securing Data and Using Encryption**
113 Understanding Permissions
113 Setting POSIX Permissions
114 Viewing POSIX Permissions
114 Interpreting POSIX Permissions
116 Modifying POSIX Permissions
116 Setting File and Folder Flags
116 Viewing Flags
116 Modifying Flags
117 Setting ACL Permissions
117 Setting ACL Permissions Using Workgroup Manager
118 Setting ACL Permissions for a File
119 Setting Global File Permissions
120 Securing Your Home Folder
120 Encrypting Home Folders
121 Using FileVault Master Keychain
122 Centrally Managing FileVault
123 Encrypting Portable Files
123 Creating a New Encrypted Disk Image
125 Creating an Encrypted Disk Image from Existing Data
125 Securely Erasing Data
126 Using Disk Utility to Securely Erase a Disk or Partition
127 Using Command-Line Tools to Securely Erase Files or Folders
127 Using Secure Empty Trash
128 Using Disk Utility to Securely Erase Free Space
128 Using Command-Line Tools to Securely Erase Free Space

Chapter 7 129 **Securing Accounts, Share Points, and Network Views**
129 Open Directory and Active Directory
130 Configuring Share Points
131 Configuring Workgroup Manager for Working with Share Points
131 Disabling Share Points
131 Restricting Access to a Share Point
133 Configuring AFP Share Points
133 Configuring SMB/CIFS Share Points
133 Configuring NFS Share Points
134 Configuring FTP Share Points

135 Controlling Network Views
136 Securing Accounts
136 Configuring User Accounts
138 Configuring Group Accounts
139 Configuring Computer Lists
140 Managing Preferences
141 Understanding Managed Preference Interaction
142 Choosing How to Manage Preferences
143 Setting the Permanence of Management
145 Managing Applications Preferences
146 Managing Classic Preferences
147 Managing Dock Preferences
149 Managing Energy Saver Preferences
150 Managing Finder Preferences
152 Managing Internet Preferences
155 Managing Login Preferences
159 Managing Media Access Preferences
161 Managing Mobility Preferences
163 Managing Network Preferences
165 Managing Printing Preferences
167 Managing Software Update Preferences
168 Managing System Preferences Preferences
169 Disabling Widgets
169 Managing Universal Access Preferences

Chapter 8
171 **Managing Certificates**
171 Understanding Public Key Infrastructure
172 Public and Private Keys
173 Certificates
173 Certificate Authorities
173 Identities
173 Self-Signed Certificates
174 Readying Certificates
174 Using Certificate Manager
175 Requesting a Certificate from a CA
176 Creating a Self-Signed Certificate
176 Importing a Certificate
177 Modifying Certificates
177 Editing a Certificate
178 Deleting a Certificate
178 Creating a Certificate Authority
178 Using Certificate Assistant
178 Creating a CA Using Certificate Assistant

180 Creating a CA from the Command Line
180 Signing a Newly Created CA
181 Storing the CA Private Key
181 Creating Folders and Files for SSL
182 Deploying Server Certificates to Clients

Chapter 9 183 **Setting General Protocols and Access to Services**
183 Setting General Protocols
184 Disabling NTP and SNMP
185 Enabling SSH
186 Setting the Server's Host Name
187 Setting the Date and Time
187 Setting Up Certificates
188 Setting Service Access Privileges

Chapter 10 191 **Securing Remote Access Services**
191 Securing Remote Login
192 Configuring Secure Shell
192 Modifying the SSH Configuration File
194 Generating Key Pairs for Key-Based SSH Connections
195 Updating SSH Key Fingerprints
196 Controlling Access to SSH
196 Understanding SSH Man-in-the-Middle Attacks
197 Transferring Files Using SFTP
198 Securing VPN Service
198 Enabling Layer Two Tunneling Protocol, Secure Internet Protocol
200 Enabling and Configuring Point-to-Point Tunneling Protocol
201 Authentication Methods
201 Offering SecurID Authentication with VPN Service
202 Configuring Access Warning Banners
203 Securing Apple Remote Desktop
203 Encrypting Observe and Control Network Data
204 Encrypting Network Data During File Copy and Package Installations
204 Securing Remote Apple Events

Chapter 11 205 **Securing Network and Host Access Services**
205 Using IPv6 Protocol
206 IPv6-Enabled Services
206 Securing DHCP Service
206 Disabling Unnecessary DHCP Services
207 Configuring DHCP Services
207 Assigning Static IP Addresses Using DHCP
208 Securing DNS Service

208 Understanding BIND
209 Turning Off Zone Transfers and Recursive DNS Queries
209 Disabling Recursion
210 Understanding DNS Security
210 DNS Spoofing
211 Server Mining
211 DNS Service Profiling
212 Denial of Service
212 ARP spoofing
212 Service Piggybacking
213 Securing Firewall Service
214 Planning Firewall Setup
214 Starting the Firewall Service
215 Creating an IP Address Group
216 Creating Firewall Service Rules
217 Creating Advanced Firewall Rules
218 Enabling Stealth Mode
219 Setting Up Firewall Service Logging
220 Securing NAT Service
220 Configuring NAT Service
221 Configuring Port Forwarding
222 Securing Bonjour Service

Chapter 12 223 **Securing Collaboration Services**
223 Disabling iChat Service
224 Securely Configuring iChat Service
225 Viewing iChat Service Logs

Chapter 13 227 **Securing Mail Service**
227 Disabling Mail Service
228 Configuring Mail Service for SSL
230 Configuring Authentication Support
232 Restricting SMTP Relay
232 Enabling Mail Filtering
233 Enabling Virus Filtering
234 Disabling the SMTP Banner

Chapter 14 235 **Securing File Services**
235 Disabling File Services
236 Choosing a File Sharing Protocol
237 Configuring AFP File Sharing Service
238 Configuring FTP File Sharing Service
240 Configuring NFS File Sharing Service

241 Configuring Windows File Sharing Service
242 Restricting File Permissions

Chapter 15 243 **Securing Web Service**
243 Disabling the Web Service
244 Disabling Web Modules
244 Disabling Web Options
245 Configuring Web Service for SSL
247 Using a Passphrase with SSL Certificates
247 Setting Up the SSL Log for a Website
248 Securing WebDAV
249 Setting Access for Websites
250 Securing Weblogs
250 Enabling Weblogs
252 Securing the Application Server
253 Securely Configuring the Application Server
253 Backing Up and Restoring Application Server Configurations
253 Securing WebObjects

Chapter 16 255 **Securing Client Configuration Management Services**
255 Securing NetBoot Service
255 Disabling NetBoot Service
256 Securely Configuring NetBoot Service
257 Viewing NetBoot Service Logs
258 Securing Software Update Service
258 Disabling Software Update Service

Chapter 17 259 **Securing Directory Services**
260 Understanding Open Directory Server Roles
260 Configuring the Open Directory Services Role
261 Starting Kerberos After Setting Up an Open Directory Master
262 Configuring Open Directory for SSL
263 Configuring Open Directory Policies
263 Setting the Global Password Policy
264 Setting a Binding Policy for an Open Directory Master and Replicas
265 Setting a Security Policy for an Open Directory Master and Replicas

Chapter 18 267 **Securing Print Service**
267 Disabling Print Service
268 Configuring Print Queues
269 Configuring Print Banners
269 Creating Banner Pages

Chapter 19 271 **Securing Multimedia Services**
271 Disabling QuickTime Streaming Server
272 Configuring a Streaming Server
272 Controlling Access to Streamed Media
273 Creating an Access File
274 Setting Up Relay Streams

Chapter 20 277 **Securing Grid and Cluster Computing Services**
277 Disabling Xgrid Service
278 Understanding Xgrid Service
278 Authenticating for the Grid
278 Single Sign-On
279 Password-Based Authentication
279 No Authentication
279 Setting Passwords for Xgrid Service
280 Securely Configuring Xgrid Service
280 Configuring an Xgrid Agent
281 Configuring an Xgrid Controller

Chapter 21 283 **Validating System Integrity**
283 Using Activity Analysis Tools
283 Configuring System Auditing
284 Installing Auditing Tools
284 Enabling Auditing
285 Setting Audit Mechanisms
285 Using the audit Tool
286 Using the auditreduce Tool
287 Using the praudit Tool
288 Deleting Audit Records
288 Audit Control Files
289 Managing Audit Log Files
289 Configuring Log Files
290 Configuring the syslogd Daemon
290 Local System Logging
291 Remote System Logging
292 Viewing Logs in Server Admin
292 About File Integrity Checking Tools
293 About Antivirus Tools

Appendix A 295 **Understanding Passwords and Authentication**
295 Understanding Password Types
295 Authentication and Authorization
296 Open Directory Passwords

296 Shadow Passwords
297 Crypt Passwords
298 Offline Attacks on Passwords
298 Password Guidelines
298 Creating Complex Passwords
299 Using an Algorithm to Create a Complex Password
300 Safely Storing Your Password
300 Password Maintenance
301 Authentication Services
301 Determining Which Authentication Option to Use
302 Password Policies
303 Single Sign-On Authentication
303 Kerberos Authentication

Appendix B 305 **Security Checklist**
305 Installation Action Items
306 Hardware and Core Mac OS X Action Items
307 Account Configuration Action Items
308 System Software Action Items
308 .Mac Preferences Action Items
308 Accounts Preferences Action Items
308 Appearance Preferences Action Items
309 Bluetooth Preferences Action Items
309 CDs & DVDs Preferences Actions Items
309 Classic Preferences Action Items
310 Dashboard and Exposé Preferences Action Items
310 Date & Time Preferences Action Items
310 Desktop & Screen Saver Preferences Action Items
310 Displays Preferences Action Items
310 Dock Preferences Action Items
311 Energy Saver Preferences Action Items
311 Keyboard and Mouse Preferences Action Items
311 Network Preferences Action Items
312 Print & Fax Preferences Action Items
312 QuickTime Preferences Action Items
312 Security Preferences Action Items
312 Sharing Preferences Action Items
313 Software Update Preferences Action Items
313 Sound Preferences Action Items
313 Speech Preferences Action Items
313 Spotlight Preferences Action Items
313 Startup Disk Preferences Action Items
314 Data Maintenance and Encryption Action Items

314 Account Policies Action Items
314 Share Points Action Items
314 Network Views Action Items
315 Account Configuration Action Items
315 Applications Preferences Action Items
315 Classic Preferences Action Items
316 Dock Preferences Action Items
316 Energy Saver Preferences Action Items
316 Finder Preferences Action Items
317 Internet Preferences Action Items
317 Login Preferences Action Items
318 Media Access Preferences Action Items
318 Mobility Preferences Action Items
318 Network Preferences Action Items
319 Printing Preferences Action Items
319 Software Update Preferences Action Items
319 System Preferences Preferences Action Items
320 Universal Access Preferences Action Items
320 Certificates Action Items
320 General Protocols and Service Access Action Items
321 Remote Access Services Action Items
322 Network and Host Access Services Action Items
322 IPv6 Protocol Action Items
322 DHCP Service Action Items
322 DNS Service Action Items
323 Firewall Service Action Items
323 NAT Service Action Items
323 Bonjour Service Action Items
323 Collaboration Services Action Items
324 Mail Service Action Items
324 File Services Action Items
325 AFP File Sharing Service Action Items
325 FTP File Sharing Service Action Items
326 NFS File Sharing Service Action Items
326 SMB/CIFS Action Items
326 Web Service Action Items
327 Client Configuration Management Services Action Items
327 Directory Services Action Items
328 Print Service Action Items
328 Multimedia Services Action Items
328 Grid and Cluster Computing Services Action Items
329 Validating System Integrity Action Items

Glossary 331

Index 343

About This Guide

Use this guide as an overview of Mac OS X Server security features that can enhance security on your computer.

This guide gives instructions for securing Mac OS X Server version 10.4 or later, and for securely managing servers and clients in a networked environment. It also provides information about the many different roles Mac OS X Server can assume in a network.

Target Audience

Administrators of server computers running Mac OS X Server version 10.4 or later are the intended audience for this guide. If you're using this guide, you should be an experienced Mac OS X Server user, be familiar with the Workgroup Manager and Server Admin applications, and have at least some experience using the Terminal application's command-line interface. You should also have experience administering a network, be familiar with basic networking concepts, and be familiar with the Mac OS X Server administration guides.

Some instructions in this guide are complex, and deviation from them could result in serious adverse effects on the server and its security. These instructions should only be used by experienced Mac OS X Server administrators, and should be followed by thorough testing.

What's New in Version 10.4

Mac OS X Server version 10.4 offers major security enhancements in the following key areas:

- **Access control lists.** Provide flexible file system permissions that are fully compatible with Windows Server 2003 Active Directory environments and Windows XP clients.
- **Secure instant messaging.** Your private, secure iChat Server, based on Jabber XMPP protocol, integrates with Open Directory for user accounts and authentication.
- **Software update server.** By enabling the new Apple Software Update Server, administrators can control which updates their users can access and when.

- **Certificate management.** Certificate Assistant is an easy-to-use application that helps you request, issue, and manage certificates.
- **Smart cards as keychains.** Use a smart card to authenticate to your computer or keychain.
- **Secure erase.** Secure erase follows the U.S. Department of Defense standard for the sanitation of magnetic media.
- **VPN service is now Kerberized.** Use Kerberos-based authentication for single sign-on to a VPN network.
- **Firewall enhanced.** The firewall service has been enhanced to use the reliable open source IPFW2 software.
- **Antivirus and antispam.** New adaptive junk mail filtering using SpamAssassin and virus detection and quarantine using ClamAV.

What's in This Guide

This guide explains how to secure servers and securely manage server and client computers in a networked environment. It does not provide information about securing clients. For help with securing computers running Mac OS X version 10.4 or later, see *Mac OS X Security Configuration*.

This guide cannot cover all possible network configurations in which Mac OS X Server might be used. Good network security and design must be used for this information to be effective, and anyone using this guide needs to be familiar with UNIX security basics, such as setting file permissions.

This guide includes the following chapters, arranged in the order that you're likely to need them when securely configuring a server.

- Chapter 1, "Introducing Mac OS X Server Security Architecture," provides an overview of the security architecture and features of Mac OS X Server. This chapter describes the security framework, access permissions, built-in security services, and directory services.
- Chapter 2, "Installing Mac OS X Server," describes how to securely install Mac OS X Server locally or remotely. This chapter also includes information about updating system software, repairing disk permissions, and securely erasing data.
- Chapter 3, "Protecting Hardware and Securing Global System Settings," describes how to physically protect your hardware from attacks. This chapter also tells you how to secure settings that affect all users of the computer.
- Chapter 4, "Securing Local Server Accounts," describes the types of user accounts and how to securely configure an account. This includes securing the accounts with strong authentication.

- Chapter 5, "Securing System Preferences," helps you configure your local server accounts securely. This includes the secure configuration of local system preferences, setting up strong authentication, credential storage, and securing data.
- Chapter 6, "Securing Data and Using Encryption," describes how to encrypt your data and how to use secure erase to ensure old data is completely removed.
- Chapter 7, "Securing Accounts, Share Points, and Network Views," describes security settings related to managed user and group accounts. This chapter also helps you set policies and enforce them using Workgroup Manager.
- Chapter 8, "Managing Certificates,"describes how to generate, request, and deploy certificates.
- Chapter 9, "Setting General Protocols and Access to Services," helps you configure general network management protocols and restrict access to other services.
- Chapter 10, "Securing Remote Access Services," tells you how to create remote connections to your server using encryption.
- Chapter 11, "Securing Network and Host Access Services," explains how to connect client computers and configure a firewall.
- Chapter 12, "Securing Collaboration Services," describes how to securely configure an iChat server so that users can communicate by chatting.
- Chapter 13, "Securing Mail Service," explains how to set up mail service to use encryption and filter for spam and viruses.
- Chapter 14, "Securing File Services," details how to configure the file services to enable secure data sharing.
- Chapter 15, "Securing Web Service," describes how to set up a web server and secure different web settings and components.
- Chapter 16, "Securing Client Configuration Management Services," tells you how to configure NetBoot securely to provide images to clients.
- Chapter 17, "Securing Directory Services," explains how to configure Open Directory service roles and password policies.
- Chapter 18, "Securing Print Service," explains how to set up print queues and banner pages.
- Chapter 19, "Securing Multimedia Services," provides security information to configure a streaming server.
- Chapter 20, "Securing Grid and Cluster Computing Services," explains how to securely configure an Xgrid agent and controller.
- Chapter 21, "Validating System Integrity," describes how to use security audits and logging to validate the integrity of your server and data.
- Appendix A, "Understanding Passwords and Authentication," describes Open Directory authentication, shadow and crypt passwords, Kerberos, LDAP bind, and single sign-on.

- Appendix B, "Security Checklist," provides a checklist that guides you through securing your server.
- The Glossary defines terms you'll encounter as you read this guide.

Note: Because Apple frequently releases new versions and updates to its software, images shown in this book might be different from what you see on your screen.

Using This Guide

The following list contains suggestions for using this guide:

- Read the guide in its entirety. Subsequent sections might build on information and recommendations discussed in prior sections.
- The instructions in this guide should always be tested in a nonoperational environment before deployment. This nonoperational environment should simulate, as much as possible, the environment where the computer will be deployed.
- This information is intended for computers running Mac OS X Server. Before securely configuring a server, determine what function that particular server will perform and apply security configurations where applicable.
- Use the security checklist in Appendix B to track and record each security task and note what settings you changed. This information can be helpful when developing a security standard within your organization.

Important: Any deviation from this guide should be evaluated to determine what security risks it might introduce. Take measures to monitor or mitigate those risks.

Using Onscreen Help

You can view instructions and other useful information from documents in the server suite by using onscreen help.

On a computer running Mac OS X Server, you can access onscreen help after opening Workgroup Manager or Server Admin. From the Help menu, select one of the following options:

- *Workgroup Manager Help* or *Server Admin Help* displays information about the application.
- *Mac OS X Server Help* displays the main server help page, from which you can search or browse for server information.
- *Documentation* takes you to www.apple.com/server/documentation, where you can download server documentation.

You can also access onscreen help from the Finder or other applications on a server or on an administrator computer. (An administrator computer is a Mac OS X computer with server administration software installed on it.) Use the Help menu to open Help Viewer, and then choose Library > Mac OS X Server Help.

To see the latest server help topics, make sure the server is connected to the Internet while you're using Help Viewer. Help Viewer automatically retrieves and caches the latest help topics from the Internet. When not connected to the Internet, Help Viewer displays cached help topics.

The Mac OS X Server Suite

The Mac OS X Server documentation includes a suite of guides that explain the services and provide instructions for configuring, managing, and troubleshooting the services. All of the guides are available in PDF format from:

www.apple.com/server/documentation/

This guide...	tells you how to:
Getting Started, Getting Started Supplement, and *Mac OS X Server Worksheet*	Install Mac OS X Server and set it up for the first time.
Collaboration Services Administration	Set up and manage weblog, chat, and other services that facilitate interactions among users.
Command-line Administration	Use commands and configuration files to perform server administration tasks in a UNIX command shell.
Deploying Mac OS X Computers for K-12 Education	Configure and deploy Mac OS X Server and a set of Mac OS X computers for use by K-12 staff, teachers, and students.
Deploying Mac OS X Server for High Performance Computing	Set up and manage Mac OS X Server and Apple cluster computers to speed up processing of complex computations.
File Services Administration	Share selected server volumes or folders among server clients using these protocols: AFP, NFS, FTP, and SMB/CIFS.
High Availability Administration	Manage IP failover, link aggregation, load balancing, and other hardware and software configurations to ensure high availability of Mac OS X Server services.
Java Application Server Guide	Configure and administer a JBoss application server on Mac OS X Server.
Mac OS X Security Configuration	Secure Mac OS X client computers.
Mac OS X Server Security Configuration	Secure Mac OS X Server computers.
Mail Service Administration	Set up, configure, and administer mail services on the server.
Migrating to Mac OS X Server From Windows NT	Move accounts, shared folders, and services from Windows NT servers to Mac OS X Server.

This guide...	tells you how to:
Network Services Administration	Set up, configure, and administer DHCP, DNS, VPN, NTP, IP firewall, and NAT services on the server.
Open Directory Administration	Manage directory and authentication services.
Print Service Administration	Host shared printers and manage their associated queues and print jobs.
QuickTime Streaming Server 5.5 Administration	Set up and manage QuickTime streaming services.
System Imaging and Software Update Administration	Use NetBoot and Network Install to create disk images from which Macintosh computers can start up over the network. Set up a software update server for updating client computers over the network.
Upgrading And Migrating	Use data and service settings that are currently being used on earlier versions of the server software.
User Management	Create and manage user accounts, groups, and computer lists. Set up managed preferences for Mac OS X clients.
Web Technologies Administration	Set up and manage a web server, including WebDAV, WebMail, and web modules.
Windows Services Administration	Set up and manage services including PDC, BDC, file, and print for Windows computer users.
Xgrid Administration	Manage computational Xserve clusters using the Xgrid application.
Mac OS X Server Glossary	Learn about terms used for server and storage products.

Getting Documentation Updates

Periodically, Apple posts revised guides and help topics. The new help topics include updates to the guides.

- To view new onscreen help topics, make sure your server or administrator computer is connected to the Internet and click the Late-Breaking News link on the main Mac OS X Server help page.
- To download the latest guides and solution papers in PDF format, go to the Mac OS X Server documentation webpage: www.apple.com/server/documentation/.

Getting Additional Information

For more information, Apple provides the following resources:

- *Read Me documents*—Important updates and special information. Look for them on the server installation discs.
- *Mac OS X Server website* (www.apple.com/macosx/server/)—Gateway to extensive product and technology information.
- *Apple Support website* (www.apple.com/support/)—Access to hundreds of articles from Apple's support organization.
- *Apple Customer Training website* (train.apple.com)—Instructor-led and self-paced courses for honing your server administration skills.
- *Apple Certification Programs website* (train.apple.com/certification/)—In-depth certification programs designed to create a high level of competency among Macintosh service technicians, help desk personnel, technical coordinators, system administrators, and other professional users.
- *Apple Discussions website* (discussions.info.apple.com)—Discussion forums for sharing questions, knowledge, and advice with other administrators.
- *Apple Product Security Mailing Lists website* (lists.apple.com/mailman/listinfo/security-announce/)—Mailing lists for communicating by email with other administrators about security notifications and announcements.
- *Open Source website* (developer.apple.com/darwin/)—Access to Darwin open source code, developer information, and FAQs.
- *Apple Product Security website* (www.apple.com/support/security/)—Access to security information and resources, including security updates and notifications.

For additional security-specific information, consult these resources:

- *NSA security configuration guides* (www.nsa.gov/snac/)—The National Security Agency provides a wealth of information about securely configuring proprietary and open source software.
- *NIST Security Configuration Checklists Repository* (checklists.nist.gov/repository/category.html)—The National Institute of Standards and Technology repository for security configuration checklists.
- *DISA Security Technical Implementation Guide* (www.disa.mil/gs/dsn/policies.html)—The Defense Information Systems Agency guide for implementing secure government networks. A Department of Defense (DoD) PKI Certificate is required to access this information.
- *CIS Benchmark and Scoring Tool* (www.cisecurity.org/bench_osx.html)—The Center for Internet Security benchmark and scoring tool used to establish CIS benchmarks.

Acknowledgments

Apple would like to thank the National Security Agency for their assistance in creating and editing the security configuration guides for Mac OS X 10.4 'Tiger' client and server.

Introducing Mac OS X Server Security Architecture

1

Mac OS X Server delivers the highest level of security through the adoption of industry standards, open software development, and smart architectural decisions.

With Mac OS X Server, a security strategy is implemented that is central to the design of the operating system, ensuring that your Mac is safe and secure. This chapter describes the features in Mac OS X Server that you can use to enhance security on your computer.

- **Open source foundation.** Using open source methodology makes Mac OS X Server a more robust, secure operating system because its core components have been subjected to peer review for decades. Apple and the larger open source community can quickly identify and fix problems.
- **Secure default settings.** When you take your Mac out of the box, it is securely configured to meet the needs of most common usage environments, so you don't have to be a security expert to set up your system. These default settings make it very difficult for malicious software to infect your system. However, security can be further configured on the computer to meet organizational or user requirements.
- **Modern security architecture.** Mac OS X Server includes state-of-the-art, standards-based technologies enabling Apple and third-party developers to build secure software for the Mac. These technologies support all aspects of system, data, and networking security required by today's applications.
- **Innovative security applications.** Mac OS X Server includes features that take the worry out of using a computer. For example, FileVault protects your documents using strong encryption, an integrated VPN client gives you secure access to networks over the Internet, and a powerful firewall secures your home network.
- **Rapid response.** Because the security of your system is so important, Apple responds rapidly to provide patches and updates. Apple works with worldwide partners, including the Computer Emergency Response Team (CERT), to notify users of any potential threats. Should vulnerabilities be discovered, the built-in software update tool automatically notifies users of security updates, which are available for easy retrieval and installation.

Security Architecture Overview

Mac OS X Server security services are built on two open source standards: Berkeley Software Distribution (BSD) and Common Data Security Architecture (CDSA). BSD is a form of the UNIX operating system that provides fundamental services and the Mac OS X file system. CDSA provides a much wider array of security services, including finer-grained access permissions, authentication of users' identities, encryption, and secure data storage. The default security settings on your Mac OS X Server computer are configured to be secure from local network and Internet attacks.

UNIX Infrastructure

The Mac OS X kernel—the heart of the operating system—is built from BSD and Mach. Among other things, BSD provides basic file system and networking services and implements a user and group identification scheme. BSD enforces access restrictions to files and system resources based on user and group IDs. Mach provides memory management, thread control, hardware abstraction, and interprocess communication. Mach enforces access by controlling which tasks can send a message to a given Mach port (a Mach port represents a task or some other resource). BSD security policies and Mach access permissions constitute an essential part of security in Mac OS X Server, and are both critical to enforcing local security.

Access Permissions

An important aspect of computer security is the granting or denying of access permissions (sometimes called access rights). A permission is the ability to perform a specific operation, such as gaining access to data or executing code. Permissions are granted at the level of folders, subfolders, files, or applications. Permissions are also granted for specific data within files or application functions.

Permissions in Mac OS X Server are controlled at many levels, from the Mach and BSD components of the kernel through higher levels of the operating system, and—for networked applications—through the networking protocols.

Security Framework

Apple built the foundation of Mac OS X Server and many of its integrated services with open source software—such as FreeBSD, Apache, and Kerberos, among many others—that has been made secure through years of public scrutiny by developers and security experts around the world. Strong security is a benefit of open source software because anyone can freely inspect the source code, identify theoretical vulnerabilities, and take steps to strengthen the software. Apple actively participates with the open source community by routinely releasing updates of Mac OS X Server that are subject to independent developers' ongoing review—and by incorporating improvements. An open source software development approach provides the transparency necessary to ensure that Mac OS X Server is truly secure.

This open approach has clear advantages and a long, well-documented history of quickly identifying and correcting source code that could potentially contain exploitable vulnerabilities. Mac OS X Server users can comfortably rely on the ongoing public examination by large numbers of security experts, which is made possible by Apple's open approach to software development. The result is an operating system that is inherently more secure.

Layered Security Defense

Mac OS X Server security is built on a layered defense for maximum protection. Security features provide solutions for securing data at all levels, from the operating system and applications to networks and the Internet.

- Secure worldwide communication—Firewall and mail filtering help prevent malicious software from compromising your computer.
- Secure applications—Authentication using keychains and encryption using FileVault helps prevent intruders from using your applications and viewing data on your computer.
- Secure network protocols—Secure sockets layer helps prevent intruders from viewing information exchange across a network and Kerberos secures the authentication process.
- Operating system—POSIX and ACL permissions help prevent intruders from accessing your files.
- Hardware—The Open Firmware Password application helps prevent people who can access your hardware from gaining root-level access to your computer files.

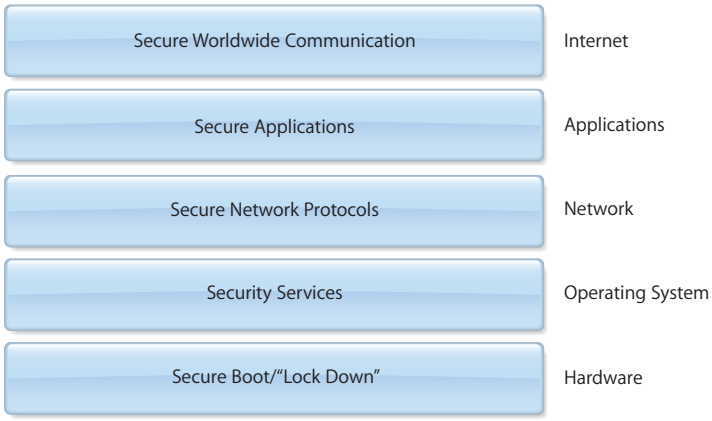

Built-In Security Services

Mac OS X Server has several security services that are managed by the security server daemon. Security server implements several security protocols such as encryption, decryption, and authorization computation. The use of the security server to perform actions with cryptographic keys enables the security implementation to maintain the keys in a separate address space from the client application, keeping them more secure.

Keychain Services

A keychain is used to store passwords, keys, certificates, and other secrets. Due to the sensitive nature of this information, the keychains use cryptography to encrypt and decrypt secrets, and they safely store secrets and related data in files.

The Mac OS X Server keychain services enable you to create keychains and provide secure storage of keychain items. Once a keychain is created, you can add, delete, and edit keychain items, such as passwords, keys, certificates, and notes for one or more users. A user can unlock a keychain with a single password and applications can then use that keychain to store and retrieve data, such as passwords.

Secure Transport Services

Secure Transport is used to implement Secure Sockets Layer (SSL) and Transport Layer Security (TLS) protocols. These protocols provide secure communications over a TCP/IP connection, such as the Internet, by using encryption and certification exchange.

Certificate, Key, and Trust Services

The certificate, key, and trust services include functions to:

- Create, manage, and read certificates
- Add certificates to a keychain
- Create encryption keys
- Manage trust policies

These functions are carried out when the services call a variety of Common Security Service Manager (CSSM) functions. This is all transparent to users.

Authorization Services

Authorization services give applications control over access to specific operations within an application. For example, a directory application accessible to any user may use authorization services to restrict access for modifying directory items to administrators. In contrast, BSD provides access permissions only to an entire file or application.

Smart Card Services

A smart card is a plastic card (similar in size to a credit card) or a USB dongle that has memory and a microprocessor embedded in it. The smart card is capable of both storing information and processing it. Smart cards can securely store passwords, certificates, and keys. A smart card normally requires a personal identification number (PIN) or biometric measurement (such as a fingerprint) as an additional security measure. Because it contains a microprocessor, a smart card can carry out its own authentication evaluation offline before releasing information. Smart cards can exchange information with a computer through a smart card reader.

Directory Services

Directory services, a critical component of any modern network environment, let you centralize information about users, groups, and computing resources in your organization. Maintaining this data in a central repository makes it possible for all servers on the network to access the same user accounts, settings, and authentication services. Directory services improve the security and manageability of your network environment, thereby reducing administration costs.

Mac OS X Server uses Open Directory as its directory service, making it easy to integrate Mac OS X client and server computers with your existing network infrastructure. The standards-based architecture provides compatibility with other LDAP servers and even with environments that use proprietary services, such as Microsoft's Active Directory and Novell's eDirectory. And for organizations that haven't yet deployed centralized directory services, the Open Directory server in Mac OS X Server offers an easy-to-deploy solution that scales to meet the needs of virtually any network environment.

Open Directory Authentication Architecture

The Open Directory authentication architecture stores password enforcement policies and authentication credentials in a robust, central repository. By assigning parameters to passwords, such as password length, types of characters needed, and expiration time, administrators can require users to pick more secure passwords.

Open Directory integrates MIT's open source Kerberos Key Distribution Center (KDC) for secure access to network resources. This robust directory-based authentication mechanism enables single sign-on to all authorized systems and services. Instead of authenticating to each service individually, you type in your password once at login to prove your identity to the Kerberos authentication authority. In response, the KDC issues strongly encrypted electronic *tickets,* used to assure all participating applications and services that you are securely authenticated. Kerberized applications and services include:

- Safari
- Telnet
- SSH
- SMB/CIFS
- Mail
- VPN
- Apple Filing Protocol

Mac OS X Server lets users share Windows-managed networks, with a home folder on either a Mac or a PC. Network administrators can set one authentication policy for all users, permitting Mac OS X users to log in and authenticate to Microsoft's proprietary Active Directory—without any specific changes to accommodate Mac OS X users.

Mac OS X Server supports Microsoft's NTLM version 2 authentication protocol for increased compatibility.

Policy Management

Using Server Admin, you can configure authentication methods for users whose password type is Open Directory (the most secure password type). Policies can be implemented to enforce network rules. This includes changing passwords at next login, granting or denying access to services or resources, and managing the methods used for authentication.

Authorization versus Authentication

Authorization is the process by which an entity, such as a user or a computer, obtains the right to perform a restricted operation. Authorization can also refer to the right itself, as in "Anne has the authorization to run that program." Authorization usually involves first authenticating the entity and then determining whether it has the appropriate permissions.

Authentication is the process of verifying the identity of a user or service. Authentication is normally done as a step in the authorization process. Some applications and operating system components carry out their own authentication. Authentication might use authorization services when necessary.

Network Deployment Considerations

Careful planning incorporating security concerns must precede deployment of Mac OS X Server in any network architecture. The server administration guides provide worksheets to assist in this process. Providing adequate isolation of the site network from the outside world, and properly separating functions for the computers within the site network, are basic security goals in designing a network.

Network Isolation

The site's connection to external networks, such as the Internet, must be properly protected. In general, this involves using a firewall to filter network traffic. The firewall should prevent unwanted access to your network and its resources from computers on the external network. For example, it's common to set up file-sharing services, such as AFP or SMB on a local network. Such services should not be available to external users, and certainly not to external networks or the Internet at large. A properly configured firewall can prevent external users from accessing the file server.

Other measures, such as intrusion detection systems, proxy servers, and host-based firewalls, can further bolster network defenses.

- For more information about designing a network's external connections, see *Firewalls and Internet Security—Repelling the Wily Hacker, 2nd Edition*, by William Cheswick and Steven Bellovin (Addison-Wesley Professional, 2003).
- For information about configuring network boundary devices and using them as firewalls, see *Router Security Configuration Guide*, located on the web at www.nsa.gov/snac/.

Functional Separation

Any computer on a local area network can be classified into one of three main categories: directory servers, other servers, and clients. Each computer on the network should fall into just one of these categories.

Directory servers are distinguished from other types of servers because they are used to manage user and client settings and contain user authentication data. Planning the structure of the hierarchy of directory servers, including replicas and backups, is especially important to ensure availability to all users. The "Open Directory Planning" chapter in the Open Directory administration guide provides a detailed explanation of this planning process.

Directory servers should be kept in a physically secure location to which nonadministrative personnel do not have access, and network access to these servers should be as restricted as possible. Only administrator users should be able to log directly in to a directory server. Examples of directory services are: Apple's LDAP-based Open Directory Server included with Mac OS X Server, Microsoft's Active Directory, and Sun's NIS/NIS+.

A typical network also includes servers for network services, such as email, file sharing, logging, and web. To the extent possible, each network service should be hosted on a separate server. Physical access should be restricted to administrative personnel wherever possible, network access should be restricted to only that which is operationally necessary, and only administrators should be able to log directly in to a server.

Client computers provide user access to the network, but do not provide any services to the rest of the network. Security-relevant settings on the client should be enforced to the extent possible. Configuration guides from NSA exist for Mac OS X, Solaris, and Microsoft Windows. The Center for Internet Security publishes configuration guidance for computers running Linux, FreeBSD, and HP-UX.

Using Accounts Securely

Many security features are built into Mac OS X Server, including industry-standard digital signatures, encryption for Apple's Mail application, and authentication for the Safari web browser. Other security features include:

- **Security system preferences.** Lets the user configure FileVault and control some aspects of authorization on the computer.
- **FileVault.** Uses 128-bit AES encryption to encrypt everything in the user's home folder. As long as the user is authenticated and logged in, the system automatically decrypts any file the user opens. However, no other user can gain access to these files.
- **Accounts system preferences.** When a user installs Mac OS X Server on a computer, that user automatically becomes a member of the admin group on the computer. Subsequently, the user or any other member of the admin group can use Accounts system preferences to add new users to the computer. For each new user, the administrator can specify whether that user is a member of the admin group.
- **Keychain Access application.** Keychain Access gives users access to Keychain Services. A user can see the passwords, certificates, and other data that are stored in their keychain. They can create new keychains, add and delete keychain items, lock and unlock keychains, and select one keychain to be the default.

Installing Mac OS X Server

2

Though the default installation of Mac OS X Server is highly secure, it can be customized for your particular network security needs.

By securely configuring the different stages of the installation process and understanding Mac OS X Server permissions, you can make sure that your computer is hardened to match your security policy.

Important: Computers should remain isolated from the operational network until they are completely and securely configured whenever possible. You should use an isolated test network for installation and configuration.

System Installation Overview

Although secure configuration of an existing Mac OS X Server installation is possible, securely configuring a fresh installation is much simpler. This might not always be practical, but it is the recommended way to configure Mac OS X Server.

The preinstallation of Mac OS X Server on a new computer is not "locked down" from a security standpoint. This is by design, since a server is used to administer an entire network and typically needs more services available. If a previous installation of Mac OS X Server exists on a computer, consider a clean installation of Mac OS X Server by doing an Erase and Install or reformatting the volume.

WARNING: Erase and Install will completely erase the content of the chosen volume. Be sure to back up your files before continuing.

When backing up and restoring any information, use the following guidelines:

- Only user files and data should be saved and later restored. Restoring system settings might change the system configuration.
- Reinstall applications from the original media. Do not restore them from a backup.

When you configure your new partitions, you should securely erase the partition that you're installing Mac OS X Server on. For more information, see "Using Disk Utility to Securely Erase a Disk or Partition" on page 126.

If you decide against securely erasing the partition, securely erase free space after installing Mac OS X Server. For more information, see "Using Disk Utility to Securely Erase Free Space" on page 128.

There are several ways to install the operating system depending on your environment and installation strategy. These include:

- Installing locally from CD or DVD
- Installing remotely from discs or images
- Installing from the command line
- Automating installation

Disabling the Open Firmware Password

If the Open Firmware password was previously enabled, disable it before beginning installation. When you set a firmware password, it prevents others from starting up the computer from a volume other than the one you have chosen as the startup disk (chosen in the Startup Disk preference panel within System Preferences). Once security is enabled, you cannot start up from other devices, such as an external FireWire disk, a CD or DVD drive, or another partition or disk inside the computer.

To disable the Open Firmware password:

1. Restart the computer while holding down the Command, Option, O, and F keys.
2. Enter the Open Firmware password when prompted.

   If you are not prompted to enter a password, the Open Firmware password is already disabled.
3. Enter the following commands:

```
$ reset-nvram
$ reset-all
```

You can also use the Open Firmware Password application to disable an Open Firmware password. For Mac OS X Server 10.4 or later, you must use the updated version provided on the software installation disc (located in the /Applications/Utilities/ folder on the disc).

For more information, see "Configuring Open Firmware Settings" on page 56.

Installing Locally from CD or DVD

The easiest and most secure way to install Mac OS X Server is to install it physically at the computer, known as a local installation, using the CD or DVD. When performing a local installation, it is recommended, if applicable, that the entire drive be reformatted using at least a 7-pass secure erase, rather than just the partition where Mac OS X Server is to be installed.

If your server has multiple partitions and you are only installing on a partition, you should still do a secure erase of that partition. This type of installation is known as a clean installation. After the entire drive is securely erased and formatted, partitions can be created as required. When installing the Mac OS X Server, only install the packages that are needed. All data on the target drive will be lost during the installation process.

> *WARNING:* The following instructions will cause all information about the target volume (disk or partition) to be lost. Back up any data on the volume that should be retained.

To install server software locally:

1. If you'll be performing the recommended clean installation rather than upgrading, preserve any user data that resides on the disk or partition you'll install the server software onto.
2. Turn on the computer and insert the first Mac OS X Server installation disc into the optical drive.
3. Restart the computer while holding down the C key. The computer starts up from the installation disc. You can release the C key when you see the Apple logo.
4. After the computer restarts, choose the language you want the server to use and click Continue.
5. When Installer opens, choose Utilities > Open Disk Utility to securely erase the target disk before proceeding.

   To securely erase and format the entire disk or partition, use Disk Utility. For more information, see "Using Disk Utility to Securely Erase a Disk or Partition" on page 126.
6. Proceed through the panes of the Installer by following the onscreen instructions.
7. When the Select a Destination pane appears, select a target disk or volume (partition) and make sure it's in the expected state.
8. At the "Install" screen, click Customize.

9 Deselect any options not needed on this server.

Any unneeded languages should not be installed.

Note: By default, X11 is not selected. The X11 X Window system provides the ability to run X11-based applications under Mac OS X Server. While this capability can be useful, it introduces configuration and security issues.

10 Click the Printer Drivers disclosure triangle to reveal individual printer drivers, and deselect any drivers that will not be needed. (Printer drivers can always be installed at a later date if a new printer is added.)

Only drivers for the printers that will be used should be installed.

11 Continue the installation. During installation, progress information is displayed. Insert the next installation disc, if prompted.

12 Verify the media by allowing the check to run. Do not skip.

After the installation is complete, the server restarts automatically and you can perform initial server setup. For further installation procedures, see "Initializing Server Setup" on page 39.

Important: Don't install additional software or store user data on the hard disk or partition where the operating system is installed. By separating the user data from Mac OS X Server, you lessen the risk of losing data if you ever reinstall or upgrade system software. If you must store additional software or data on the system software partition, consider mirroring the drive.

Installing Remotely from Disks or Images

An alternative method for installing Mac OS X Server is to install it from a separate computer. This process is known as remote installation. When retrieving the image over a network, make sure that the network is isolated and can be trusted. When doing a remote server installations, you must know the target server's identity and password.

- **The identity of the target server.** When using Server Assistant, you must be able to recognize the target server in a list of servers on your local subnet or enter the IP address of the server if it resides on a different subnet. Information provided for servers in the list include IP address, host name, and MAC address (also called hardware or Ethernet address). The IP address is assigned by a DHCP server on the network. If no DHCP server exists, the target server uses a 169.254.xxx.xxx address unique among servers on the local subnet. Later, when you set up the server, you can change the IP address.
- **The preset password for the target server.** The password consists of the first eight digits of the built-in hardware serial number of the server. To find the serial number of the server, look for a label on the server. Older computers don't have built-in hardware serial numbers. For these computers, use 12345678.

If this installation is being performed on an existing server, and you know the identity and an administrator name and password of the server, you can obtain the serial number using command-line tools.

To obtain the serial number using command-line tools:

1 Open Terminal.

2 Log in to the target server using `ssh`. For information about `ssh`, see "Configuring Secure Shell" on page 192 or the `ssh` man page.

3 Use the `ioreg` command to get the serial number.

```
$ ioreg -l | grep SerialNumber | head -n 1 | awk '{print $4}'
```

Installing Remotely from the Installation Discs

To install Mac OS X Server on a remote server from the server installation discs, you need an administrator computer to manage the installation and access to the target computer. Since you are installing remotely, you will need someone physically at the computer to start it up using the installation discs.

To install to a remote server using the installation discs:

1 If you'll be performing a clean installation rather than upgrading, preserve any user data that resides on the disk or partition you'll install the server software on.

Important: Don't install additional software or store user data on the hard disk or partition where the operating system is installed. With this approach, you won't risk losing those files if you must reinstall or upgrade system software. If you must install additional software or store data on the system software partition, consider mirroring the drive.

2 Open Disk Utility to partition and format the disk as Mac OS Extended.

3 Start the target computer from the first installation disc. The procedure you use depends on the target server hardware.

- If the target server has a keyboard and an optical drive, insert the first installation disc into the optical drive. Then hold down the C key while restarting the computer.
- If the target server is an Xserve with a built-in optical drive, start the server using the first installation disc by following the instructions in the Xserve user's guide for starting from an installation.
- If the target server is an Xserve with no built-in optical drive, you can start it in target disk mode and insert the installation disc into the optical drive on your administrator computer. You can also use an external FireWire optical drive or an optical drive from another Xserve to start the server from the installation disc. Instructions for using target disk mode and external optical drives are in the quick start guide or Xserve user's guide that came with your Xserve.

4 On an administrator computer, navigate to /Applications/Server/ and open Server Assistant (you don't have to be an administrator on the local computer to use Server Assistant). Select "Install software on a remote server."

5 Identify the target server.

If it's on the local subnet, select it in the list. Otherwise, click "Server at IP Address" and enter an IP address in IPv4 format (xxx.xxx.xxx.xxx).

6 When prompted for a password, type the first eight digits of the built-in hardware serial number of the server. To find a serial number of a server, look for a label on the server. If you're installing on an older computer that has no built-in hardware serial number, use 12345678 for the password.

7 Proceed by following the onscreen instructions.

8 When the "Volumes" pane appears, select a target disk or volume (partition) and make sure it's in the expected state. Then select it and click Continue.

Important: When you perform an upgrade installation, make sure saved setup data won't be inadvertently detected and used by the server. If saved setup data is used, existing server settings will be overwritten by the saved settings. For more information about saved setup data, see the getting started guide.

9 During installation, progress information is displayed. Insert the next installation disc, if prompted.

- While installation proceeds, you can open another Server Assistant window to install server software on another computer; choose File > New Window to do so.
- After installation is complete, the target server restarts automatically and you can perform initial server setup. For further installation procedures, see "Initializing Server Setup" on page 39.

Installing Remotely from an Image

If you need to install server software on a large number of servers or if you need to reinstall server software frequently, you can automate installation by using an installation image that resides on disk rather than on the installation discs.

There are two types of disk images that you can use to accomplish remote installation:

- A boot image is a file that looks and acts like a mountable disk or volume and contains the system software needed to act as a startup disk for computers through the network. You can install Boot images on a computer, eliminating the need to use CD or DVD discs.
- An install image is a special boot image that starts up the client long enough to install software from the image, after which the computer can start up from its own hard drive.

You can install and reinstall the server software with a known good and secure configuration that you previously set up on a model server using disk images. Before creating the disk image, follow the procedures on "Installing Locally from CD or DVD" on page 33 and "Initializing Server Setup" on page 39 to ensure a secure configuration.

Before creating the disk image, you might also want to securely configure additional server settings as described in the other chapters of this guide. Make sure the model server you are imaging meets all of the security requirements of your organization and is thoroughly tested.

For information about how to create and install server software with disk images, see the getting started guide.

Secure erase can be done before installation or after installation. It's best to do it before installation to ensure that old data is completely overwritten and not recoverable. Before installing server software from a disk image, securely erase the physical disk or partition that the image is being installed on using at least a 7-pass erase. For more information, see "Securely Erasing Data" on page 125.

To install server software using a disk image:

1. Open System Image Utility.
2. Create a Network Install image from the server installation CDs or DVD or from a "model" version 10.4 server you've already set up.
3. Use Server Admin to start NetBoot service and enable the disk image.
4. Start each target computer so it starts up using the disk image.
5. Open Server Assistant and choose "Install software on a remote server," and then identify the computers that started from the network image when you get to the Destination pane.
6. Proceed as you would to install server software on any remote computer.

   The system imaging and software update administration guide describes how to create and deploy disk images.

Installing from the Command Line

There are two command-line tools you can use to install Mac OS X Server, `installer` and `asr`. These tools can be used locally or remotely.

Using the installer Command Tool for Installation

Use the `installer` tool to install server software on a local or remote computer from the command line. For detailed information about `installer`:

- See the command-line administration guide.
- Open the Terminal application and type `installer`, `installer -help`, or `man installer`.

For information about using `installer` to install server software, see the getting started guide. If you follow the instructions for performing a clean installation, back up the user files you want to preserve, then use `diskutil` to securely erase (7-pass or 35-pass) the volume and format it to enable journaling.

To secure erase a volume with 7-pass erase:

```
$ diskutil secureErase "Case-sensitive HFS+" 2 "/Volumes/Mount 01"
```

For more information, see "Securely Erasing Data" on page 125. You can also use `diskutil` to partition the volume and to set up mirroring. For more information, see the `diskutil` man page.

Important: Don't store data on the hard disk or hard disk partition where the operating system is installed. With this approach, you won't risk losing data should you need to reinstall or upgrade system software. If you must store additional software or data on the system software partition, consider mirroring the drive.

Using the asr Command Tool for Installation

You can use Apple Software Restore (ASR) to restore from a network-based disk image created using the Network Install pane of System Image Utility or by using Disk Utility. With ASR, you can restore an image deployed by an ASR server or you can save that image to a disk. By saving the image to disk, you can verify its validity by using Disk Utility to checksum the image before using it. ASR can be much more efficient than using NetBoot service to deploy disk images, especially when refreshing computers simultaneously.

You can configure ASR to continually send out a stream of networking data over the network. This is called multicast ASR. Multiple computers can connect to this stream of data simultaneously. Computers can connect to the same stream of data at any time. Since all computers are refreshed using the same stream of data, and not a separate stream for each computer, the server and the network are not as heavily strained as when deploying with NetBoot service.

It is possible to overload the network when using a multicast ASR server, reducing available bandwidth for other services. Improperly configuring the ASR data rate option can create a denial of service situation. `asr` is a command-line tool. For more information about `asr` options, see the `asr` man page.

To image a volume:

```
$ sudo asr -source /Volumes/Classic -target /Volumes/install
```

> ***WARNING:*** When restoring a system image onto a volume, the target drive will be erased.

To restore a system image onto a volume:

```
$ sudo asr -source compressedimage -target <targetvol> -erase
```

You can use `asr` to multicast images to several computers at once. However, you must be sure that you are on a trusted network, including all reachable hosts and subnets.

See the `asr` man page for more information.

Initializing Server Setup

After installing Mac OS X Server, the computer restarts and loads Server Assistant which you can use to interactively configure the server, locally or remotely.

Using Server Assistant

Server Assistant eases the process of configuring your server. Server Assistant will configure the administrator account, network settings and services, date, and time. Before you begin configuring your server, fill out the Mac OS X Server worksheet included in the appendix of the getting started guide. This worksheet helps you gather all of the pertinent target server information before proceeding.

Note: After installation of server software is complete, the computer restarts and Server Assistant opens automatically. If you want to postpone server setup until a later time, press Command-Q. The computer shuts down. When it's restarted, Server Assistant opens automatically.

Server Assistant can run locally from the server, or remotely from an administrator computer that can connect to the target server.

It is important to avoid storing the password used to access and configure a server in a keychain on a general purpose administrator computer which might have other users accessing it. This would increase the risk of an unauthorized user gaining access to servers. Do not store the server password in any keychain to prevent potential leakage.

It is a good practice when installing Mac OS X Server locally to complete the configuration prior to connecting the server to the network.

Setting Up a Secure Local Server

After server software has been installed on a server, you can set the server up locally if you have physical access to the computer.

To set up and secure a server interactively:

1. When the server is restarted, Server Assistant opens automatically.
2. Enter the setup data you've recorded on the worksheet as you move through the Assistant's panes, following the onscreen instructions.
3. In the Keyboard pane, choose your keyboard layout for the server, and click Continue.
4. In the Serial Number pane, enter the Mac OS X Server serial number along with the Site License information.
5. Click Continue to display the Administrator Account pane. In this step, you create a system administrator account. The password settings will also be used for the root account. As such, you should take special care to ensure that this account is as secure as possible.
   a. Limit the number of administrator accounts issues. This makes it easier to retain control over the computer and identify whether a particular activity noted in the logs was legitimate.
   b. When entering the administrator account information for both the Name and the Short Name fields, use names other than "administrator," "admin," or some form of the word administrator. The name alone should not identify the account as an administrator account.
   c. Use a strong password in the Password and Verify fields. Passwords can be up to 255 characters long and contain uppercase letters, lowercase letters, numbers, and special characters. Choose a password that consists of at least 12 characters that would not be found in a dictionary, and that contains mixed-case letters, numbers, and special characters.
   d. After setting up the administrator user, click Continue.
6. In the Network Names pane, you identify the server for accessing from the network. When entering the Computer Name and Local Hostname, the names should not indicate the purpose of the computer. The word "server" should not be used as the name or part of the name. Click Continue.
7. From the Network Interfaces pane, select only those interfaces that will be used and deselect all others. For example, if the network interface for the server will be Built-in Ethernet only, deselect Built-in FireWire. AppleTalk should not be used. Click Continue.
8. In the TCP/IP Connection pane, select "Manually" for the Configure IPv4 setting. The use of DHCP or BootP is not recommended. Make sure that any DHCP or DNS servers you specify for the server you're setting up to use are running. Click Continue.
9. In the Directory Usage pane, the "Set directory usage" setting should be set to Standalone Server to simplify the installation process. The type of directory usage depends on the role of the server being installed. See Chapter 17, "Securing Directory Services," on page 259 for information about configuring directory usage. Click Continue.

10 In the Services pane, do not enable any services yet. The services that should be enabled depend on the role of the server being installed. Each service should be configured carefully before activation. Click Continue.

11 In the Time Zone pane, select the time zone that the server is located in. Click Continue.

12 When setting your server time on the Network Time pane, a network timeserver should be specified if a local timeserver is available. Either select not to use a network timeserver, or select the "Use a network timeserver" box, and type the name or address of the local timeserver in the NTP Server field.

   Some authentication services, including Kerberos, require that time be synchronized across all computers, which necessitates synchronization with a timeserver. For security, one timeserver on the local network can synchronize with a trusted Internet timeserver, but it is the only server that should do so. Direct use of an Internet timeserver is not recommended for other servers. Click Continue.

   Note: If NTP is to be used on a network without Internet access, the server providing the NTP service needs to either have another time source connected, such as a GPS unit, or needs to be set up to use an undisciplined local clock. See www.ntp.org.

13 In the Date & Time pane, set the server date and time. Click Continue.

14 After all setup data has been entered, Server Assistant displays a summary of the data in the Confirm Settings pane. Review the setup data you entered. Optionally, click Go Back to change it.

   To save the setup data as a text file or in a form you can use for automatic server setup (a saved setup file or saved directory record), click Save As.

   To encrypt the file or record, select "Save in Encrypted Format" and then enter and verify a passphrase. You must supply the passphrase before an encrypted setup file can be used by a target server.

15 To initiate setup of the local server, click Apply.

   When server setup is complete, click Restart Now. Now you can log in as the administrator user created during setup to configure services.

Setting Up a Secure Remote Server

After server software has been installed on a server, you can use the interactive approach to set it up remotely from an administrator computer that can connect to the target server.

To set up and secure a remote server interactively:

1 Make sure the target server is running.

2 On an administrator computer, open Server Assistant from /Applications/Server/. You don't have to be an administrator to use Server Assistant.

3 In the Welcome pane, select "Set up a remote server" and click Continue.

4 In the Destination pane, put a check in the Apply column for the remote server you want to set up, then type its preset password in the Password field and click Continue to connect to the server.

   If you don't see the target server on the list, click Add to add it or Refresh to determine whether it's available.

5 In the Language pane, specify the language you want to use to administer the target server.

6 Click Continue and enter the setup data as you would if the server was local. For details, see "Setting Up a Secure Local Server" on page 39.

   When server setup is complete, click Continue. The target server restarts automatically and you can log in as the administrator user created during setup to configure services.

Updating System Software

After installing Mac OS X Server, be sure to install the latest approved security updates. Mac OS X Server includes Apple Software Update, an application that downloads and installs software updates either from Apple's Software Update server or from an internal software update server. You can configure software update so that it checks for updates either periodically or whenever you choose. You can also configure software update to download, but not install, updates, in case you want to install them later.

Before installing updates, check with your organization for their policy on downloading updates. They might prefer that you use an internal software update server, which reduces the amount of external network traffic and allows the organization to prequalify software updates against organization configurations prior to updating individual computers.

Software updates should be installed immediately after the operating system installation. Software updates are obtained and installed in different ways:

- Using Apple Software Update to download and install updates from an internal software update server
- Using Apple Software Update to download and install updates from Internet-based software update servers
- Manually downloading and installing updates as separate software packages

Important: All security updates published by Apple contain fixes for security issues, and are usually released in response to a specific known security problem. Applying these updates is essential.

If Apple Software Update does not install an update that you request, contact your network administrator. Failure to update signifies that the requested update could be a malicious file.

Important: Before connecting to the Internet, you should ensure that your network services are securely configured. If you have not secured and validated your settings for network services, you should not enable your network connection to install software updates. Until you have securely configured your network services settings, you are limited to using the manual method of installing software updates. For more information, see "Securing Software Update Preferences" on page 106.

Updating from an Internal Software Update Server

The computer automatically looks for software updates from an internal software update server. By using an internal software update server, you reduce the amount of data transferred outside of the network. Your organization can control which updates can be installed on your computer.

If you run software update over a wireless or an untrusted network, you run a chance of downloading malicious updates from a rogue software update server. Software update, however, will not install a package that has not been digitally signed by Apple prior to distribution.

If you connect your computer to a network that manages its client computers, the network can require that the computer use a specified software update server.

To specify your software update server:

```
$ defaults write com.apple.SoftwareUpdate CatalogURL http://
     swupdate.apple.com:8088/index.sucatalog
```

where *swupdate.apple.com* is the fully qualified domain name (FQDN) or IP address of your software update server.

Updating from Internet-Based Software Update Servers

Software update can periodically check the Internet for software updates. Instead of using your operational computer to check for and install updates, consider using a test-bed computer to download updates and verify file integrity before installing updates. You can then transfer the update packages to your operational computer. See "Updating Manually from Installer Packages" on page 44.

To download and install software updates using Software Update:

1 Choose Apple (🍎) > Software Update.

After Apple Software Update looks for updates to your installed software, it displays a list of all updates. To get older versions of updates, go to the software update website at www.apple.com/support/downloads/.

2 Select the updates you want to install, and choose Update > Install and Keep Package. When you keep the package, it is stored in the /Library/Packages/ folder. If you do not want to install any of the updates, click Quit.

3 Accept the licensing agreements to start installation.

Some updates might require your computer to restart. If, after installing updates, software update asks you if you want to restart the computer, do so.

Important: Make sure updates are installed when the computer can be restarted without affecting the users accessing the server.

Updating Manually from Installer Packages

Software updates can be manually downloaded for all of Apple's products from www.apple.com/support/downloads/ using a computer designated specifically for downloading and verifying updates. The download should be done separately so that file integrity can be verified before the updates are installed.

It is possible to review the contents of each security update before installing it. To see the contents of a security update, go to Apple's Security Support Page at www.apple.com/support/security and click the "Security Updates page" link.

To manually download, verify, and install software updates:

1 Go to www.apple.com/support/downloads/ and download the necessary software updates on a computer designated for verifying software updates.

Note: Updates provided through Apple Software Update might sometimes appear earlier than the standalone updates.

2 Review the SHA-1 digest (also known as checksum) for each update file downloaded, which should be posted online with the update package.

3 Check all downloaded updates for viruses.

4 Verify the integrity of each update.

For more information, see "Verifying the Integrity of Software" on page 45.

5 Transfer the downloaded update packages from your test computer to your current computer. The default download location for update packages is /Library/Packages/. You can transfer update packages to any location on your computer.

6 Double-click the package. If the package is located within a disk image (dmg) file, double-click the dmg file, and then double-click the package.

7 Proceed through the installation steps.

8 Restart the computer, if requested.

Install the appropriate software update and then install any subsequent security updates. These updates should be installed in order by release date, oldest to newest.

Verifying the Integrity of Software

Software images and updates can include an SHA-1 digest, which is also known as a checksum. You can use this SHA-1 to verify the integrity of the software. Software updates retrieved and installed automatically from Apple Software Update verify the checksum before installation.

To verify software integrity:

1. Open Terminal.
2. Use the `sha1` command to display a file's SHA-1 digest.

   ```
   $ /usr/bin/openssl sha1 full_path_filename
   ```

   The *`full_path_filename`* is the full path filename of the update package or image for which the SHA-1 digest is being checked.

   If provided, the SHA-1 digest for each software update or image should match the digest created for that file. If it does not, the file was corrupted in some way and a new copy should be obtained.

Repairing Disk Permission

Permissions on files can sometimes be set incorrectly, especially during a software installation. Incorrect permissions can cause the computer to malfunction and even introduce security vulnerabilities. Repairing these permissions is recommended after performing any software installation on Mac OS X Server.

Important: The procedure for repairing disk permissions should be performed after every software installation, including the operating system, updates, and applications.

Kinds of Permissions

Before you change or repair disk permissions, you should understand the two kinds of file and folder permissions that Mac OS X Server supports:

- Portable Operating System Interface (POSIX) permissions—standard for UNIX operating systems.
- Access Control Lists (ACLs) permissions—used by Mac OS X Server, are compatible with Microsoft Windows Server 2003 and Microsoft Windows XP.

Note: In this guide, the term "privileges" refers to the combination of ownership and permissions, while the term "permissions" refers only to the permission settings that each user category can have (Read & Write, Read Only, Write Only, and None).

POSIX Permissions Overview

POSIX permissions let you control access to files and folders. Every file or folder has read, write, and execute permission defined for three different categories of users (owner, group, and everyone). There are four types of standard POSIX access permissions that you can assign: Read & Write, Read Only, Write Only, and None.

For more information, see "Setting POSIX Permissions" on page 113.

ACL Permissions Overview

Access Control List provides an extended set of permissions for a file or folder and enables you to set multiple users and groups as owners. An ACL is a list of access control entries (ACEs), each specifying the permissions to be granted or denied to a group or user, and how these permissions are propagated throughout a folder hierarchy. In addition, ACLs are compatible with Windows Server 2003 and Windows XP, giving you added flexibility in a multiplatform environment.

ACLs provide more granularity when assigning privileges than POSIX permissions. For example, rather than giving a user full writing permissions, you can restrict him or her to the creation of only folders and not files.

If a file or folder has no ACEs defined for it, Mac OS X Server applies the standard POSIX permissions. If a file or folder has one or more ACE defined for it, Mac OS X Server starts with the first ACE in the ACL and works its way down the list until the requested permission is satisfied or denied. After evaluating the ACEs, Mac OS X Server evaluates the standard POSIX permissions defined for the file or folder. Then, based on the evaluation of ACL and standard POSIX permissions, Mac OS X Server determines what type of access a user has to a shared file or folder.

For more information, see "Setting ACL Permissions" on page 117.

Using Disk Utility to Repair Disk Permissions

Installing software sometime causes file permissions to become incorrectly set. Incorrect file permissions can create security vulnerabilities. Disk Utility only repairs POSIX permissions or the minimal ACL permissions.

Most software you install in Mac OS X Server is installed from package (.pkg) files. Each time something is installed from a package file, a Bill of Materials (.bom) file is stored in the packages receipt file. Each Bill of Materials file contains a list of the files installed by that package, along with the proper permissions for each file.

When you use Disk Utility to verify or repair disk permissions, it reads the Bill of Materials files from the initial Mac OS X Server installation and compares its list to the actual permissions on each file listed. If the permissions differ, Disk Utility can repair them.

You should repair disk permissions if you are experiencing symptoms that would indicate permission related problems after installing software, software updates, or applications.

Note: If you've modified permissions for files in accordance with organizational policies, be aware that repairing disk permissions can reset those modified permissions to those stated in the Bill of Materials files. After repairing permissions, you should reapply the file permission modifications to stay within your organizational policies.

To repair disk permissions:

1. Open Disk Utility.
2. Select the partition that you want to repair.

   Ensure that you select a partition, not a drive. Partitions are contained within drives and are indented one level in the list on the left.
3. Click Repair Disk Permissions.

   If you do not select a partition, this button is disabled.
4. Choose Disk Utility > Quit Disk Utility.
5. Choose Installer > Quit Installer, and click Restart.

Protecting Hardware and Securing Global System Settings

3

After installing and setting up Mac OS X Server, protect your hardware and secure global system settings.

Protecting hardware serves as a first line of defense from unauthorized access. Another line of defense involves securing global system settings. Securing global system settings involves removing Mac OS 9 and locking down Mac OS X Server startup, system swap space, and fast user switching. You can also configure log files for monitoring computer activity.

Once your global system settings are securely configured, a login window banner can be set up to warn unauthorized users of the consequences of misuse.

Protecting Hardware

The most secure method for physically protecting your server hardware is protection from unwanted physical access. If someone can physically access a computer, it becomes much easier to compromise the computer's security. When someone has physical access to the computer, they can install malicious software or various event-tracking and data-capturing services.

Use as many layers of physical protection as possible. Restrict access to rooms that contain computers that store or access sensitive information. Provide room access only to individuals who must use those computers. If possible, lock the computer in a secure container when it is not in use, or bolt or fasten it to a wall or piece of furniture.

The hard drive is the most critical hardware component in your computer. Take special care to prevent access to the hard drive. If someone removes your hard drive and installs it in another computer, they can bypass any safeguards you set up. Lock or secure the computer's internal hardware. If you can't guarantee the physical security of the hard drive, consider using FileVault for each home folder (FileVault encrypts home folder content and prevents the content from being compromised).

If you have a portable computer, keep it secure. Lock up the computer or hide it when it is not in use. When transporting the computer, never leave it in an insecure location. Consider buying a computer bag with a locking mechanism, and lock the computer in the bag when you aren't using it.

A computer left unattended and logged in can be a security risk. To protect your computers from being used when on and unattended, you should enable a password protected screen saver. See "Securing Security Preferences" on page 104.

Disabling Hardware

Hardware components such as wireless features and microphones should be physically disabled if possible. Only an Apple Certified Technician should physically disable these components, which may not be practical in all circumstances. The following instructions provide an alternative means of disabling these components by removing the associated kernel extensions. Removing the kernel extensions does not permanently disable the components; however, administrative access is needed to restore and reload them. Although disabling hardware in this manner is not as secure as physically disabling hardware, it is more secure than only disabling hardware through the System Preferences. This method of disabling hardware components may not be sufficient to meet site security policy. Consult operational policy to determine if this method is adequate.

The following instructions will remove AirPort, Bluetooth, the microphone, and support for an external iSight camera. This will not remove the support for the internal iSight cameras currently shipping on some Macintosh systems. There is currently no way to disable this camera in software without disabling all USB drivers, which will also disable the keyboard, mouse, etc.

Important: Repeat these instructions every time a system update is installed.

To remove kernel extensions for certain hardware:

1. Open the /System/Library/Extensions folder.
2. To remove AirPort support, drag the following files to the Trash:

   AppleAirPort.kext

   AppleAirPort2.kext

   AppleAirPortFW.kext
3. To remove support for Bluetooth, drag the following files to the Trash:

   IOBluetoothFamily.kext

   IOBluetoothHIDDriver.kext

4 To remove support for audio components such as the microphone, drag the following files to the Trash:

   AppleOnboardAudio.kext

   AppleUSBAudio.kext

   AudioDeviceTreeUpdater.kext

   IOAudioFamily.kext

   VirtualAudioDriver.kext

5 To remove support for the iSight camera, drag the following file to the Trash:

   Apple_iSight.kext

6 (Optional) To remove support for mass storage devices (e.g. USB flash drives, external USB hard drives, external FireWire Hard Drives), drag the following files to the Trash:

   IOUSBMassStorageClass.kext

   IOFireWireSerialBusProtocolTransport.kext

7 Open the /System/Library folder.

8 Drag the following files to the Trash:

   Extensions.kextcache

   Extensions.mkext

9 Choose Finder > Secure Empty Trash to delete the file.

10 Restart the system.

Removing Mac OS 9

When you upgrade from previous versions of Mac OS X Server to Mac OS X Server version 10.4, an adaptation of Mac OS 9, known as "Classic," remains on the computer. If you perform a new installation of Mac OS X Server version 10.4 without upgrading, Mac OS 9 is not installed on the computer. It is possible to install Mac OS 9 on computers with a new installation of Mac OS X Server.

Mac OS 9 lacks many of the security features included with Mac OS X Server, so you should remove it unless you need it. If you need to use Mac OS 9, you can run it from a CD or DVD, or from a disc image.

Using the Command Line to Remove Mac OS 9

To remove Mac OS 9, use the command line. You must log in as an administrator who can use the `sudo` command to remove files. For more information, see "Securing the Local System Administrator Account" on page 70.

> ***WARNING:*** Incorrectly entering a folder name in the following command could result in removal of the Mac OS X Server operating system or all Mac OS X Server applications. Use caution when removing the files.

To remove Mac OS 9 and Mac OS 9 applications and files:

1. Log in to Mac OS X Server as an administrator who can use `sudo` to remove files.

   By default, all users who are administrators can use the `sudo` command to remove files. If you modify /etc/sudoers, you can choose which users can use `sudo`. For more information, see the `sudoers` man page.

2. Open Terminal.

3. Enter the following command to remove the Classic icon from System Preferences:

```
$ sudo rm -rf '/System/Library/PreferencePanes/Classic.prefPane'
```

4. Enter the following commands to remove Classic folders and files:

```
$ sudo rm -rf '/System/Library/Classic/'
$ sudo rm -rf '/System/Library/CoreServices/Classic Startup.app'
$ sudo rm -rf '/System/Library/User Template/English.lproj/Desktop/Desktop
     (Mac OS 9)'
```

5. Enter the following commands to remove Mac OS 9 folders and files:

```
$ sudo rm -rf '/System Folder'
$ sudo rm -rf '/Mac OS 9 Files/'
```

6. Enter the following command to remove Mac OS 9 applications:

```
$ sudo rm -rf '/Applications (Mac OS 9)'
```

7. Restart the computer.

Running Mac OS 9 from a CD or DVD

Classic is an environment for running Mac OS 9 applications. If you must run Mac OS 9, you can use Classic to run it from a CD or DVD. By running Mac OS 9 from a CD or DVD, you enforce read-only access.

Note: Intel-based Macintosh computers do not support the Classic environment or Mac O 9.

To run Mac OS 9 from a CD or DVD:

1. Install Classic and the software that requires Classic on a test-bed computer.
2. Burn the Mac OS 9 System Folder from the test-bed computer onto a blank CD or DVD.

   The System Folder is located at the base level of a partition. It might be named something besides "System Folder." System folders are denoted by a folder icon with a 9 superimposed on them.
3. Eject the CD or DVD from the test-bed computer and insert it into your operational computer.
4. Open Classic preferences on your operational computer.
5. Select the System Folder located on the CD or DVD in the "Select a system folder for Classic" list.
6. Click Start.

Running Mac OS 9 from a Disc Image

Classic is an environment for running Mac OS 9 applications. If you must run Mac OS 9, you can use Classic to run it from a disc image. By running Mac OS 9 from a disc image, you enforce read-only access.

Note: Intel-based Macintosh computers do not support the Classic environment or Mac O 9.

To run Mac OS 9 from a disc image:

1. Install Mac OS 9 and the software that requires Mac OS 9 on a test-bed computer.
2. On the test-bed computer, create a folder and name it Mac OS 9.
3. Copy the Mac OS 9 System folder into the Mac OS 9 folder you created in the previous step.
4. On the test-bed computer, open Disk Utility.
5. Choose File > New > Disk Image from Folder.
6. Select the Mac OS 9 folder (created in step 2), and click Image.
7. In Image Format, choose read-only.
8. In Encryption, choose none.
9. Click Save.
10. Copy the Mac OS 9 disc image to your operational computer.
11. Double-click the Mac OS 9 disc image to mount it.
12. Open Classic preferences on your operational computer.
13. Select the System Folder located on the mounted disc image in the "Select a system folder for Classic" list.
14. Click Start.

Securing System Startup

When a computer starts up, it first starts either Open Firmware or Extensible Firmware Interface (EFI). EFI is similar to Open Firmware, but it runs on Intel-based Macintosh computers. Open Firmware or EFI determines which partition or disk to load Mac OS X from. They also allow (or prevent) the user to enter single-user mode.

Single-user mode automatically logs in the user as "root." This is dangerous because root user access is the most powerful level of access, and actions performed as root are anonymous.

If you create an Open Firmware or EFI password, you disable single-user mode. The password also stops users from loading unapproved partitions or disks, and from enabling target disk mode at startup.

After creating an Open Firmware or EFI password, you must enter this password when you start the computer from an alternate disk (for situations such as hard drive failure or file system repair).

> ***WARNING:*** Open Firmware settings are critical. Take great care when changing these settings and when creating a secure Open Firmware password.

To secure system startup, perform one of the following tasks:

- Use the Open Firmware Password application to set the firmware password
- Set the Open Firmware password within Open Firmware
- Verify and set the security mode from the command line

Open Firmware password protection can be bypassed if the user changes the physical memory configuration of the machine and then resets the PRAM three times (by holding down Command, Option, P, and R keys during system startup). An Open Firmware password will provide some protection, however, it can be reset if a user has physical access to the machine and can change the physical memory configuration of the machine.

You can require a password to start single-user mode, which would further secure your computer.

For more information about Open Firmware password protection, see AppleCare Knowledge Base article #106482, "Setting up Open Firmware Password protection in Mac OS X 10.1 or later" (www.apple.com/support/), and AppleCare Knowledge Base article #107666, "Open Firmware: Password Not Recognized when it Contains the Letter 'U'" (www.apple.com/support/).

Using the Open Firmware Password Application

The Mac OS X Server installation disc includes the Open Firmware Password application, which allows you to enable an Open Firmware or EFI password.

> ***WARNING:*** Attempts to use firmware in a manner that is not explicitly endorsed by Apple might damage your computer's logic board. Any repairs that are necessary because of this damage will not be covered under the terms of the Apple One-Year Limited Warranty, AppleCare Protection Plan, or other AppleCare agreement.

When you set an Open Firmware password, it prevents others from starting up the computer from a volume other than the one you have chosen as the startup disk in the Startup Disk preference panel within the System Preferences. Once security is enabled, you cannot start up from other devices, such as an external FireWire disk, a CD-ROM drive, or another partition or disk inside the computer.

When an Open Firmware password is set, it:

- Blocks the ability to use the C key to start up from an optical disc.
- Blocks the ability to use the N key to start up from a NetBoot server.
- Blocks the ability to use the T key to start up in Target Disk Mode (on computers that offer this feature).
- Blocks the ability to start up in Verbose mode by pressing the Command-V key combination during startup (specific to PowerPC-based computers).
- Blocks the ability to start up in single-user mode by pressing and holding the Command-S key combination during startup (specific to PowerPC-based computers).
- Blocks a reset of Parameter RAM (PRAM) by pressing and holding the Command-Option-P-R key combination during startup.
- Requires the password to use the Startup Manager, accessed by pressing and holding the Option key during startup.
- Requires the password to enter commands after starting up in Open Firmware, which is done by pressing and holding the Command-Option-O-F key combination during startup (specific to PowerPC-based computers).
- Blocks the ability to use the D key to start up from the Diagnostic volume of the Install DVD (specific to Intel-based computers).

Important: Open Firmware password protection does not prevent someone with physical access to the computer from restarting it or turning it off. Open Firmware password protection can only effectively protect a computer that enjoys some degree of physical security.

If you reset the PRAM or Open Firmware, you must reselect your startup device prior to resetting the Open Firmware password. The Open Firmware password can be reset and changed by any one of the following:

- Any administrator user, as designated in Accounts preferences (or in Server Admin)
- Physical access to the inside of the computer
- When the computer is started up in Mac OS 9

To use the Open Firmware Password application:

1. Log in with an administrator account and open Open Firmware Password (located on the Mac OS X Server installation disc in /Applications/Utilities/).
2. Click Change.
3. Select "Require password to change Open Firmware settings."

   To disable the Open Firmware or EFI password, deselect "Require password to change Open Firmware settings." You won't have to enter a password and verify it. Disabling the Open Firmware password is only recommended when you install Mac OS X.
4. Enter a new Open Firmware or EFI password in the Password and Verify fields. Click OK.

   This password can be up to eight characters and unique to the computer.

   Do not use the capital letter U in an Open Firmware password.

   Important: Anyone who gains root access to the computer can view the Open Firmware password. If the password is used on another computer this could compromise the password.
5. Close Open Firmware Password.

   You can test your settings by attempting to load single-user mode. Restart the computer while holding down the Command and S keys. If the login window loads, changes made by Open Firmware Password completed successfully.

Configuring Open Firmware Settings

To prevent users from obtaining root access by starting in single-user mode or starting from an alternate disk, you should alter the Open Firmware settings. You can securely configure Open Firmware settings within Open Firmware. The Open Firmware security mode should be set to `command` or `full`.

- `Command`—When the value for the `security-mode` variable is `command`, the system prompts for a password when changing Open Firmware settings.
- `Full`—When the value is set to `full` the system requires a password before restarting and after restarting.

Note: If you are using an Intel-based Macintosh computer, you cannot use the following method to change the Open Firmware password. Use Open Firmware Password instead.

To configure the Open Firmware settings from within Open Firmware:

1. Restart the computer while holding down the Command, Option, O, and F keys.

   This loads Open Firmware.

2. At the prompt, change the password.

   ```
   > password
   ```

3. Enter and verify the password to be used as the Open Firmware password.

   This password can be up to eight characters and unique to the computer.

   Do not use the capital letter U in an Open Firmware password.

   Important: Anyone who gains root access to the computer can view the Open Firmware password. If the password is used on another computer this could compromise the password.

4. Set the security mode variable by entering the following command.

   ```
   > setenv security-mode command
   ```

   In command mode, the computer will only start up from the partition selected in Startup Disk.

   If you want to set the security-mode variable to the value `full`, use the following command.

   ```
   > setenv security-mode full
   ```

   Full mode is more restrictive than command mode. After enabling full mode, all Open Firmware commands will require that you enter your Open Firmware password. This includes the `boot` command, and therefore Mac OS X Server will not start up unless you enter `boot` and authenticate with the Open Firmware password.

5. Restart the computer and enable Open Firmware settings with the following command:

   ```
   > reset-all
   ```

   The login window should appear after restarting.

   You can test your settings by attempting to load single-user mode. Restart the computer while holding down the Command and S keys. If the login window appears, your Open Firmware settings are set correctly.

> *WARNING:* Modifying critical system files can cause unexpected issues. Your modified files may also be overwritten during software updates. Make these modifications on a test computer first, and thoroughly test your changes every time you change your system configuration.

Using Command-Line Tools to Secure Startup

Open Firmware can also be configured through the command line by using the `nvram` tool. However, only the `security-mode` environment variable can be securely set. The `security-password` variable should not be set from the `nvram` tool or it will be visible when viewing the environment variable list. To set the password for Open Firmware, start the computer into Open Firmware and set the password. See "Configuring Open Firmware Settings" on page 56 for more information. The `nvram` tool requires system administrator or root access to set environment variables.

Note: If you are using an Intel-based Macintosh computer, you cannot use the following method to secure startup. Use the Open Firmware Password application instead.

To use nvram to secure startup from the command line:

1. Set the security mode by entering the following command.

```
# nvram security-mode="command"
```

   If you want to set the security mode to `full`.

```
# nvram security-mode="full"
```

2. Verify that the variable has been set. The following command displays a list of all the environment variables excluding the `security-password` variable.

```
# nvram -p
```

Requiring a Password for Single-User Mode

Additional protection can be provided in case the Open Firmware (PowerPC-based systems) or EFI (Intel-based systems) password is bypassed. By requiring entry of the root password during a single-user mode boot, the system can prevent automatic root login if the OF/EFI password is compromised.

To require entry of the root password during a single-user mode boot, the console and ttys must be marked as insecure in /etc/ttys. In fact, the system will require entry of a special root password, stored in /etc/master.passwd. If this remains unset as recommended, then it will be impossible for a user to enter the root password and complete the single-user boot, even if the Open Firmware password protection was bypassed.

To require entry of the root password for single-user mode:

1. Log in as an administrator.
2. Start the Terminal application, located in /Applications/Utilities.
3. At the prompt, enter the command:

```
$ cd /etc
```

4. To create a backup copy of /etc/ttys, enter the command:

```
$ sudo mv ttys ttys.old
```

5 To edit the ttys file as root, enter the command:

```
$ sudo pico ttys
```

6 Replace all occurrences of the word "secure" with the word "insecure" in the configuration lines of the file. Any line that does not begin with a "#" is a configuration line.

7 Exit, saving changes.

Configuring Access Warnings

You can use a login window or Terminal access warning to provide notice of a computer's ownership, to warn against unauthorized access, or to remind authorized users of their consent to monitoring.

Enabling Access Warnings for the Login Window

Before enabling an access warning, check your organization's policy for what to use as your access warning.

When a user tries to access the computer's login window (either locally or through Apple Remote Desktop), the user will see the access warning you create.

To create a login window access warning:

1 Open Terminal.

2 Change your login window access warning:

```
$ sudo defaults write /Library/Preferences/com.apple.loginwindow
     LoginwindowText "Warning Text"
```

Replace *Warning Text* with your access warning text.

Your logged-in account needs to be able to use `sudo` to perform a `defaults write`.

3 Log out to test your changes.

Your access warning text appears below the Mac OS X Server title.

Enabling Access Warnings for the Command Line

Before enabling an access warning, check your organization's policy for what to use as your access warning.

When a user opens Terminal locally or connects to the computer remotely, the user sees the access warning you create.

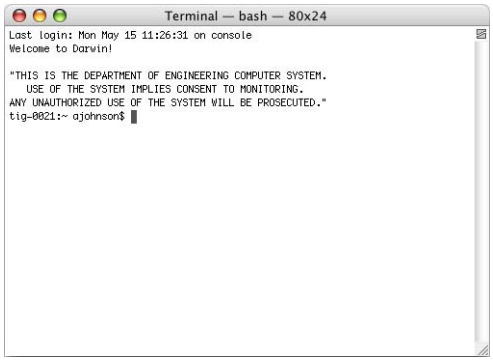

To create a command-line access warning:

1 Open Terminal.

2 Open /etc/motd in a text editor:

```
$ sudo pico /etc/motd
```

You must be able to use `sudo` to open `pico`. For information about how to use `pico`, enter `man pico` in a Terminal window.

3 Replace any existing text with your access warning text.

4 Save your changes and exit the text editor.

5 Open a new Terminal window to test your changes.

Your access warning text appears above the prompt in the new Terminal window.

Securing Fast User Switching

When using fast user switching, all users should authenticate using authentication credentials. Users should not use the automatic login feature. Using fast user login with automatic login causes a security issue because users can freely switch between users without any authentication. If you decide to use fast user switching, ensure the users on the Mac OS X Server are trusted users because any process that one user starts will run in the background while switched to another user.

Important: Fast user switching should not be considered when multiple users are accessing local accounts. For example, any local volume mounted automatically mounts with the other persons permissions on their login, giving them access to content that they most likely should never have. Fast user switching is only recommended when a single user has full control of the local accounts and the system configuration.

To enable fast user switching:

1. Choose System Preferences > Accounts.
2. Click the lock in the lower left-hand corner to authenticate using your authentication credentials.
3. Click "Login Options" and deselect "Automatically log in as."
4. Select the "Enable fast user switching" and select the "View name" that will be displayed in the top right-hand corner of the screen.
5. Click the lock again to secure the settings.

Displaying a Login Warning Banner

A login banner can be used to provide notice of the system's ownership, give legal warning to unauthorized users, and remind authorized users of their consent to monitoring. The text displayed in the login banner should be determined by site policy. Warning banners should be displayed on all systems. Banners should be provided to anyone logging onto the system.

Setting a Local Login Warning Banner

You can configure your computer to display a login warning banner when a local user logs on.

To provide a login warning banner to any local users:

1. Convert the binary property list com.apple.loginwindow.plist to an XML file using the following command:

```
$ sudo plutil -convert xml1 /Library/Preferences/com.apple.loginwindow.plist
```

2 Open the file /Library/Preferences/com.apple.loginwindow.plist as an administrator using a plain text editor:

```
$ sudo vi /Library/Preferences/com.apple.loginwindow.plist
```

3 Immediately after the `<dict>` tag, add new lines with a `<key>` and `<string>` entry, as shown in bold in the following command. The new `<key>` tag must contain `LoginwindowText`, but the new `<string>` can contain whatever warning banner has been indicated by site policy.

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE plist PUBLIC "-//Apple Computer//DTD PLIST 1.0//EN" "http://
     www.apple.com/DTDs/PropertyList-1.0.dtd">
<plist version="1.0">
<dict>
<key>LoginwindowText</key>
<string>THIS IS A DEPARTMENT OF DEFENSE COMPUTER SYSTEM. USE OF THE SYSTEM
     IMPLIES CONSENT TO MONITORING. ANY UNAUTHORIZED USE OF THE SYSTEM WILL
     BE PROSECUTED.
</string>
...
```

4 Save changes and exit.

5 Convert com.apple.loginwindow.plist XML back to a binary property list using the following command:

```
$ sudo plutil -convert binary1 /Library/Preferences/
     com.apple.loginwindow.plist
```

The warning banner should appear for the next person logging in to the computer.

Setting a Login Warning Banner for Remote Services

You can configure your computer to display a login warning banner when a remote user logs on.

To provide a login warning banner to users logging in to remote services on the system:

1 Open the /etc/motd file as an administrator.

```
$ sudo vi etc/motd
```

2 Enter the warning banner that has been approved. For example:

```
THIS IS A DEPARTMENT OF DEFENSE COMPUTER SYSTEM. USE OF THE SYSTEM IMPLIES
     CONSENT TO MONITORING. ANY UNAUTHORIZED USE OF THE SYSTEM WILL BE
     PROSECUTED.
```

3 Exit, saving changes. The warning banner should appear for the next person logging in to a remote service.

Securing Local Server Accounts

4

Securely configuring local user accounts requires determining how the accounts will be used and setting the level of access for users.

When you define a local user's account, you specify the information needed to prove the user's identity: user name, authentication method (such as a password, digital token, smart card, or biometric reader), and user identification number (user ID). Other information in a user's account is needed by various services to determine what the user is authorized to do and to personalize the user's environment.

Types of User Accounts

A user account stores data that Mac OS X Server needs to validate the user's identity and provide services for the user. When you log in to Mac OS X Server, you can use either a nonadministrator account or an administrator account. The main difference between the two types of accounts is that there are safety mechanisms to prevent nonadministratior users from editing key preferences, or from performing certain actions that are critical to computer security. Administrator users are not as limited as nonadministrator users.

The nonadministrator and administrator accounts can be further defined by specifying additional user privileges or restrictions depending on the account usage.

User Account	User Access
Standard nonadministrator	Non-privileged user access
Managed nonadministrator	Restricted user access
Server administrator	Administer the server configuration
Directory domain administrator	Administer the configured domains on the server
System administrator (root)	Unrestricted access to the entire server

Unless administrator access is required, you should always log in as a nonadministrator user. You should log out of the administrator account when you are not using the computer as an administrator. If you are logged in as an administrator, you are granted some privileges and abilities that you might not need. For example, you can change some system preferences without being required to authenticate. This automatic authentication bypasses a security safeguard that prevents malicious or accidental modification of system preferences.

Note: This chapter describes how to secure local accounts configured on Mac OS X Server. For more information about securing user and group network accounts using Workgroup Manager, see Chapter 7, "Securing Accounts, Share Points, and Network Views."

General Guidelines for Securing Accounts

When you create user accounts, follow these guidelines:

- Never create accounts that are shared by several users. Each user should have his or her own standard or managed account.

  Individual accounts are necessary to maintain accountability. System logs can track activities to each user account, but if several users share the same account, it becomes much more difficult to track which user performed a certain activity. Similarly, if several administrators share a single administrator account, it becomes much harder to track which administrator performed a specific action.

  If someone compromises a shared account, it is less likely to be noticed. Users might mistake malicious actions performed by an intruder for legitimate actions by one of the users sharing the account.

- Each user needing administrator access should have an individual administrator account in addition to a standard or managed account. Administrator users should only use their administrator accounts for administrative purposes.

  By requiring an administrator to have a personal account for typical use and an administrator account for administrator purposes, you reduce the risk of an administrator inadvertently performing actions like accidentally reconfiguring secure system preferences.

Defining User IDs

A user ID is a number that uniquely identifies a user. Mac OS X computers use the user ID to keep track of a user's folder and file ownership. When a user creates a folder or file, the user ID is stored as the creator ID. A user with that user ID has read and write permissions to the folder or file by default.

The user ID should be a unique string of digits between 500 and 2,147,483,648. New users created using the Accounts pane of System Preferences are assigned user IDs starting at 501. It is risky to assign the same user ID to different users, because two users with the same user ID have identical folder and file permissions.

The user ID 0 is reserved for the root user. User IDs below 100 are reserved for system use; users with these user IDs should not be deleted and should not be modified except to change the password of the root user. If you do not want the user to appear in the login window of computers with Mac OS X version 10.4 or later installed, assign a user ID of less than 500.

In general, once user IDs have been assigned and users start creating files and folders, you shouldn't change user IDs. One possible scenario in which you might need to change a user ID is when merging users created on different servers onto one new server or cluster of servers. The same user ID might have been associated with a different user on the previous server.

Here are the predefined user accounts:

Name	Short name	UID
Unprivileged User	nobody	-2
System Administrator	root	0
System Services	daemon	1
Printing Services	lp	26
Postfix User	postfix	27
VPN MPPE Key	vpn_*nnnnnnnnnnn*	57
World Wide Web Server	www	70
Apple Events User	eppc	71
MySQL server	mysql	74
sshd Privilege separation	sshd	75
QuickTime Streaming Server	qtss	76
Cyrus IMAP User	cyrus	77
Mailman User	mailman	78
Application Server	appserver	79
Clamav User	clamav	82
Amavisd User	amavisd	83
Jabber User	jabber	84
Xgrid Controller	xgridcontroller	85
Xgrid Agent	xgridagent	86
Application Owner	appowner	87
WindowServer	windowserver	88
Unknown User	unknown	99

Here is a list of the predefined group accounts:

Short name	Group ID
nobody	-2
nogroup	-1
wheel	0
daemon	1
kmem	2
sys	3
tty	4
operator	5
mail	6
bin	7
staff	20
lp	26
postfix	27
postdrop	28
utmp	45
uucp	66
dialer	68
network	69
www	70
mysql	74
sshd	75
qtss	76
mailman	78
appserverusr	79
admin	80
appserveradm	81
clamav	82
amavisd	83
jabber	84
xgridcontroller	85
xgridagent	86
appowner	87
windowserver	88
accessibility	90
unknown	99

Securing Local Nonadministrator Accounts

There are two types of nonadministrator accounts: standard and managed. Standard users do not have administrator privileges, and do not have any parental controls limiting their actions. Managed users also do not have administrator privileges, but they have active parental controls. Parental controls help deter unsophisticated users from performing a few malicious activities. They can also help prevent users from accidentally misusing their computer. Parental controls are a subset of the managed preferences for Workgroup Manager.

When creating nonadministrator accounts, you should restrict the accounts so that they can only use what is operationally required. For example, if you plan to store all data on your local computer, you can disable the ability to burn DVDs.

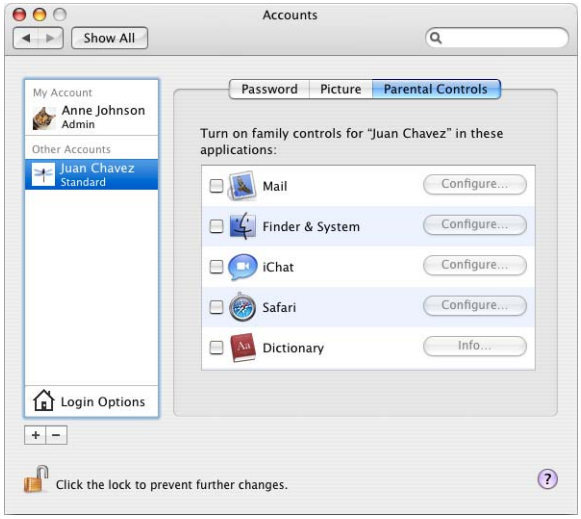

To secure a locally managed account:

1. Open Accounts preferences.
2. Click the lock to authenticate. Enter an administrator's name and password, and click OK.

   You can also authenticate through the use of a digital token, smart card, or biometric reader.
3. Select an account labeled as "Standard" or "Managed."

   You cannot set parental controls on administrator users. When selecting a user with the "Managed" label, make sure you do not select an account with preferences managed through the network.
4. Click Parental Controls.
5. Select Finder & System, and click Configure.
6. Click Some Limits.

   You can also enable Simple Finder, which restricts an account to only using applications listed in the Dock. With Simple Finder enabled, users cannot create or delete files. It also prevents users from being able to change their own passwords. Enabling Simple Finder is not recommended unless your computer is used in a kiosk-like environment.
7. Select "Open all System Preferences," and "Change password."

   To enable "Change password," you must enable "Open all System Preferences." By allowing the user to open all System Preferences, you also allow the user to change settings like the timing settings for activating the screen saver. These settings can impact security. However, the inability for a user to change his or her own password is also a security risk.
8. Deselect "Burn CDs and DVDs."
9. Deselect "Administer printers."
10. Deselect "Allow supporting programs."

    If you allow supporting programs, applications can load "helper" applications. If these helper applications are insecure, they can expose your computer to other security risks. These helper applications are loaded by an application, and not by you, so you might not be aware of them running.
11. Select "This user can only use these applications."

12 Deselect applications and utilities that are not approved for use.

When you install third-party applications, they may be added to this list. You should disable all third-party applications unless the user has a specific need to use the application, and can do so in a secure manner. Additionally, if you're connecting to an organization's network, you should only install third-party applications that are specifically approved by the organization.

13 Deselect "Applications (Mac OS 9)" and "Others."

14 Click OK.

Securing Local Server Administrator Accounts

A user account with administrator privileges can perform all standard user-level tasks and key administrator-level tasks, such as:

- Create user accounts
- Change the FileVault master password
- Enable or disable sharing
- Enable, disable, or change firewall settings
- Change other protected areas within System Preferences
- Install system software

In addition to restricting the distribution of administrator accounts, you should also limit the use of administrator accounts. Each administrator should have two accounts: a standard account for daily use, and an administrator account for when administrator access is needed.

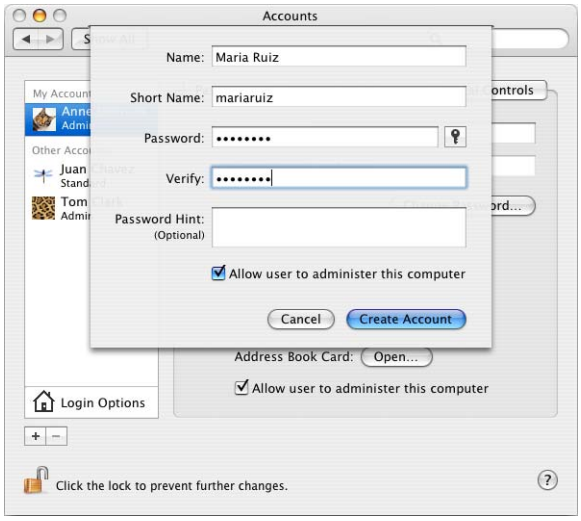

Securing a Local Directory Domain Administrator Account

A directory domain can reside on a computer running Mac OS X Server (for example, the LDAP folder of an Open Directory master, a NetInfo domain, or other read/write directory domain) or it can reside on a non-Apple server (for example, a non-Apple LDAP or Active Directory server). Only a directory domain administrator can change the directory domain, including the managed accounts in the directory domain.

When configuring a directory domain administrator account, follow the same security guidelines as you would with any other administrator account.

You can modify the /etc/authorization configuration file to change authorizations for administrators and standard users.

To modify authorization, change the /etc/authorization file:

1. Edit the /etc/authorization file using the `vi` tool, which allows for safe editing of the file. The command must be run as root:

   ```
   $sudo vi /etc/authorization
   ```

2. Enter your password when prompted.
3. This will display the property list for authorization, listing all available keys.
4. Locate the key you want to modify. For example to change who has access to unlock the screensaver, modify the `system.login.screensaver` key by changing the rule:

   ```
   <key>rule</key>
     <string>authenticate-session-owner-or-admin</string>
   ```

   to

   ```
   <key>rule</key>
     <string>authenticate-session-owner</string>
   ```

   Doing this restricts the administrator from unlocking the screensaver.

5. Save and quit `vi`.

Securing the Local System Administrator Account

The most powerful user account in Mac OS X Server is the system administrator, or root, account. The root account is primarily used for performing UNIX commands. Generally, any actions that involve critical system files require that you perform those actions as root. Even if you are logged in as an administrator, you still have to perform these commands as root, or by using the `sudo` command. Mac OS X Server logs all actions performed using the `sudo` command. This helps you track any misuse of the `sudo` command on a computer.

You can use the `su` command to log in to the command line as another user. By entering `su root`, you can log in as the root user. You can use the `sudo` command to perform commands that require root privileges. You should restrict access to the root account.

If multiple users can log in as root, it is impossible to track which user performed root actions. Direct root login should not be allowed, because the logs cannot identify which administrator logged in. Instead, accounts with administrator privileges should be used for login, and then the `sudo` command used to perform actions as root.

Note: By default, `sudo` is enabled for all administrator users. You can disable root login or restrict the use of `sudo` command.

To remove the ability of the root user to log in:

1. Open NetInfo Manager.
2. Click Security > Authentication. Authenticate when requested.
3. Click Security > Disable Root User.
4. Click Domain > Save Changes.
5. Exit NetInfo Manager.

Restricting sudo Usage

The computer uses a file named /etc/sudoers to determine which users have the authority to use `sudo`. You can modify root user access by changing the /etc/sudoers file to restrict `sudo` access to certain accounts, and allow those accounts to perform only specifically-allowed commands. This granularity gives you fine control over what users can do as root. For information about how to modify the /etc/sudoers file, see the `sudoers` man page.

The list of administrators allowed to use `sudo` should be limited to only those administrators who require the ability to run commands as root.

To restrict sudo usage:

1. Edit the /etc/sudoers file using the `visudo` tool, which allows for safe editing of the file. The command must be run as root:

   ```
   $sudo visudo
   ```

2. Enter the administrator password when prompted.

   Note: There is a time-out value associated with `sudo`. This value indicates the number of minutes until `sudo` prompts for a password again. The default value is 5, which means that after issuing the `sudo` command and entering the correct password, additional `sudo` commands can be entered for five minutes without reentering the password. This value is set in the /etc/sudoers file. For more information, see the `sudo` and `sudoers` man pages.

3. In the Defaults specification section of the file, add the following line:

   ```
   Defaults timestamp_timeout=0
   ```

4 Restrict which administrators are allowed to run `sudo` by removing the line that begins with `%admin`, and adding the following entry for each user, substituting the user's short name for the word *user*:

```
user ALL=(ALL) ALL
```

Doing this means that any time a new administrator is added to a system, that administrator must be added to the /etc/sudoers file as previously described, if that administrator requires the ability to use `sudo`.

5 Save and quit `visudo`.

For more information, see the `sudoers` man pages.

Using Strong Authentication

Authentication is the process of verifying the identity of a local or network user. Mac OS X Server supports local and network-based authentication to ensure that only users with valid authentication credentials can access the computer's data, applications, and network services.

Passwords can be required at login, to wake the system from sleep or a screen saver, to install applications, or to change system settings. Mac OS X Server also supports emerging authentication methods, such as smart cards, digital tokens, and biometric readers.

Strong authentication is created using combinations of the following three authentication dimensions:

- What the user knows, such as a password or PIN number.
- What the user has, such as a SecurID card, smart card, or drivers license.
- What the user is, such as a finger print, retina, or DNA.

Using a combination of all three dimensions above makes authentication more reliable and user identification more certain. If the use of all three dimensions is not practicable, combine as many as possible.

Using Password Assistant

Mac OS X Server includes Password Assistant, an application that analyzes the complexity of a password, or generates a complex password for you. You can specify the length and type of password you'd like to generate. For example, you can create a randomly generated password, or a FIPS-181 compliant password.

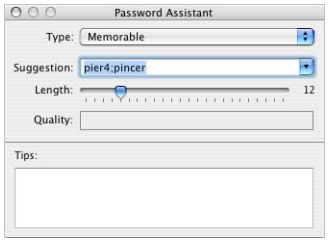

You can open Password Assistant from certain applications. For example, when you create a new account or change passwords in Accounts preferences, you can use Password Assistant to help you create a secure password.

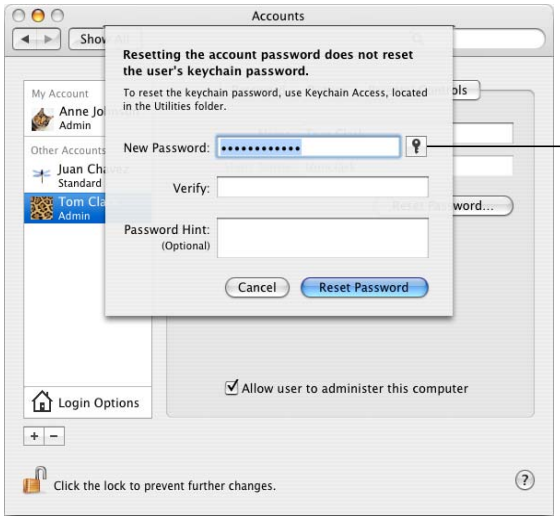

For more information, see "Creating Complex Passwords" on page 298.

Using Smart Cards

A smart card can be a plastic card (similar in size to a credit card) or a USB dongle that has memory and a microprocessor embedded in it. The smart card is capable of both storing information and processing it. Smart cards can securely store passwords, certificates, and keys. A smart card normally requires a personal identification number (PIN) or biometric measurement (such as a fingerprint) as an additional security measure. Because it contains a microprocessor, a smart card can carry out its own authentication evaluation offline before releasing information. Smart cards can exchange information with a computer through a smart card reader.

For more information, see the *Smart Card Setup Guide* located on the web at www.apple.com/itpro/federal/.

Using Tokens

A digital token is used to identify a user for commerce, communication, or access control. This token can be generated by either software or hardware. Some of the most common tokens are the RSA SecurID and the Cryptogram KT-1. These are hardware devices that automatically generate tokens to identify the user. The generated tokens are specific to that user, so two users with different RSA Steroids or different Cryptogram KT-1s will have different tokens.

You can use tokens for two-factor authentication. *Two-factor* refers to authenticating both through something you have (a One-Time-Password token) and something you know (a fixed password). The use of tokens increases the strength of the authentication process.

Tokens are frequently used for VPN authentication. For information, see "Securing VPN Service" on page 198.

Using Biometrics

Mac OS X Server supports emerging biometrics-based authentication technologies, such as thumbprint readers. Password-protected websites and applications can now be accessed without having to remember a long list of passwords. Some biometric devices allow you to authenticate simply by placing your finger on the pad. Unlike a password, your fingerprint can never be forgotten or stolen. Fingerprint identity products provide personal authentication and network access. The use of biometrics adds and additional factor to authentication by using something you are (your fingerprint).

Storing Credentials

Mac OS X Server includes Keychain Access, an application that manages collections of passwords and certificates into a single credential store called a keychain. Each keychain can hold a collection of credentials and protect them with a single password. Keychains store encrypted passwords, certificates, and any other private values (called secure notes). These values are accessible only by unlocking the keychain using the keychain password and only by applications that have been approved and added to the access control application list.

You can create multiple keychains, each of which appears in a keychain list in Keychain Access. Each keychain can store multiple values; each value is called a key item. You can create a new key item in any user-created keychain. When an application needs to store a key item in a keychain, it stores it in the one designated as your default. The default is the keychain named login, but you can change that to any user-created keychain. The default keychain is denoted by the name being displayed in bold.

Each key item on the keychain has an access control list (ACL) that can be populated with applications that have authority to use that key item. Further restriction can be added that forces an application with access to confirm the keychain password.

The main issue with having to remember many passwords is that you're likely to either make all the passwords identical or keep a written list of all passwords. By using keychains, you can greatly reduce the number of passwords that you have to remember. Since you no longer have to remember passwords for a multitude of accounts, the passwords chosen can be very complex and could even be randomly generated.

Keychains provide some additional protection for passwords, passphrases, certificates, and other credentials stored on the system. Also, in some cases, such as using a certificate to sign an email message, the certificate must be stored in a keychain. If a credential must be stored on the system, store and manage it using Keychain Access. Check your organization's policy on keychain use.

Using the Default User Keychain

When a user's account is first created, a single, default keychain named "login" is created for that user. The password for the login keychain is initially set to the user's login password and is automatically unlocked when the user logs in. It remains unlocked unless the user locks it, or until the user logs out.

The settings for the login keychain should be changed so that the user will be required to unlock the login keychain when he or she logs in, or after waking the computer from sleep.

To secure the login keychain:

1. Open Keychain Access.
2. If you do not see a list of keychains, click Show Keychains.
3. Select the login keychain.
4. Choose Edit > Change Password for Keychain "login."
5. Enter the current password, and create and verify a new password for the login keychain.

   By creating a login keychain password that is different from the normal login password, your keychain will not be automatically unlocked at login.

   You can use Password Assistant to help you create a more secure password. For information about how to use Password Assistant, see "Using Password Assistant" on page 73.
6. Choose Edit > Change Settings for Keychain "login."
7. Select "Lock when sleeping."
8. Deselect "Synchronize this keychain using .Mac."
9. Secure each individual login keychain item.

   For information, see "Securing Keychain Items" on page 77.

Securing Keychain Items

Keychains can store multiple encrypted items. You can configure some of these individual items so that only certain applications are permitted access. Access control cannot be set for certificates.

To secure individual keychain items:

1. In Keychain Access, select a keychain, and then select an item.
2. Click the Information ("i") button.
3. Click Access Control. Authenticate if you are requested to do so.
4. Select "Confirm before allowing access."

   By enabling this option, Mac OS X Server prompts you before giving a security credential to an application.

   If you selected "Allow all applications to access this item" you allow any application to access the security credential whenever the keychain is unlocked. When accessing the security credential, there is no user prompt, so enabling this is a security risk.
5. Select "Ask for Keychain password."

   After selecting this, you have to provide the keychain password before applications can access security credentials. Enabling this is particularly important for critical items, such as your personal identity (your public certificate and the corresponding private key) that is needed when signing, or decrypting information. These items can also be placed in their own keychains.
6. Remove all nontrusted applications that are listed in "Always allow access by these applications," by selecting each application and clicking the Remove (–) button.

   Any application listed here will be prompted to enter the keychain password to access the security credentials.

Creating Additional Keychains

When a user account is created, it contains only the initial default keychain, login. A user can create additional keychains, each of which can have different settings and purposes.

For example, a user might want to group all his credentials for mail accounts into one keychain. Since mail programs query the server frequently to check for new mail, it would not be practical to expect the user to reauthenticate every time such a check is being performed. The user could create a keychain and configure its settings such that he or she would be required to enter the keychain password at login and whenever the computer is awakened from sleep mode. He or she could then move all items containing credentials for mail applications into that keychain and set each item so that only the mail application associated with that particular credential can automatically access it. This would force all other applications to authenticate to access that credential.

Configuring a keychain's settings for use by mail applications might be unacceptable for other applications. If a user has an infrequently used web-based account, it would be more appropriately stored in a keychain configured to require reauthentication for every access by any application.

Additionally, you can create multiple keychains to accommodate varying degrees of sensitivity. By separating your keychains by sensitivity, you prevent the exposure of your more sensitive credentials to less sensitive applications with credentials on the same keychain.

To create a keychain and customize its authentication settings:

1. In Keychain Access, choose File > New Keychain.
2. Enter a name and select a new location for the keychain. Click Create.
3. Enter a password and verify it. Click OK.
4. If you do not see a list of keychains, click Show Keychains.
5. Select the new keychain.
6. Choose Edit > Change Settings for keychain "*keychain_name*." Authenticate, if requested.
7. Change the "Lock after # minutes of inactivity" setting depending on the access frequency of the security credentials included in the keychain.

   If the security credentials are frequently accessed, do not select "Lock after # minutes of inactivity."

   If the security credentials are moderately accessed, select "Lock after # minutes of inactivity" and select an appropriate value, such as 15. If you use a password-protected screensaver, consider setting this value to the time required for your screensaver to start.

   If the security credentials are accessed infrequently, select "Lock after # minutes of inactivity," and select an appropriate value, such as 1.
8. Select "Lock when sleeping."
9. Drag the desired security credentials from other keychains to the new keychain. Authenticate, if requested.

   You should have keychains that only contain related certificates. For example, you could have a mail keychain that only includes mail items.
10. If you are asked to confirm access to the keychain, enter the keychain password and click Allow Once.

    After confirming access, Keychain Access moves the security credential to the new keychain.
11. Secure each individual item in the security credentials for your keychain.

    For information, see "Securing Keychain Items" on page 77.

Using Portable and Network-Based Keychains

If you're using a portable computer, consider storing all of your keychains on a portable drive, such as a USB flash memory drive. The portable drive can be removed from the portable computer and stored separately when the keychains are not in use. Anyone attempting to access data on the portable computer will need the portable computer, the portable drive, and the password for the keychain stored on the portable drive. This provides an extra layer of protection if the laptop is stolen or misplaced.

To use a portable drive for storing keychains, you'll have to move all your keychain files to the portable drive, and configure Keychain Access to use the keychains on the portable drive. The default location for your keychain is ~/Library/Keychains/. However, it is possible to store keychains in other locations.

You can further protect portable keychains by storing them on biometric USB flash memory drives, or by storing your portable drive contents in an encrypted file. For information, see "Encrypting Portable Files" on page 123.

Check with your organization to see if they allow you to use portable drives to store keychains.

To set up a keychain for use from a portable drive:

1. Open Keychain Access.
2. If you do not see a list of keychains, click Show Keychains.
3. Choose Edit > Keychain List.
4. Note the location of the keychain that you want to set up. The default location is /System/Library/Keychains/. Click Cancel.
5. Select the keychain that you want set up.
6. Choose File > Delete Keychain "*keychain_name*."
7. Click Delete References.
8. Copy the keychain files from the previously noted location to the portable drive.
9. Move the keychain to the trash and use Secure Empty Trash to securely erase the keychain file stored on the computer.

   For information, see "Using Secure Empty Trash" on page 127.
10. Open Finder, and double-click the keychain file located on your portable drive to add it to your keychain.

Securing System Preferences

5

Securing Mac OS X Server system preferences enables further protection against attacks.

System Preferences has many different configurable preferences within it that can be enabled to further enhance the system security. Some of these configurations might be things to consider, depending on your organization.

Mac OS X Server includes many system preferences that you can customize to improve security. When modifying settings for one account, make sure your settings are mirrored on all other accounts, unless there is an explicit need for different settings.

You can view system preferences by choosing Apple > System Preferences. In the System Preferences window, click any of the individual preferences to view them.

Some of the more critical preferences require that you to authenticate before you can modify their settings. To authenticate, you click a lock and enter an administrator's name and password (or use a digital token, smart card, or biometric reader). If you log in as a user with administrator privileges, these preferences are unlocked. If you log in as a standard user, these preferences remain locked. After unlocking preferences, you can lock them again by clicking the lock.

Preferences that require authentication include the following:

- Accounts
- Date & Time
- Energy Saver
- Network
- Print & Fax
- Security
- Sharing
- Startup Disk

This chapter lists each set of preferences included with Mac OS X Server and describes modifications suggested to improve security.

Securing .Mac Preferences

.Mac is a suite of Internet tools designed to help you synchronize your data and other important information when you're away from the computer. You should not use .Mac if you must store critical data only on your local computer. You should only transfer data over a secure network connection to a secure internal server.

If you must use .Mac, enable it only for user accounts that don't have access to critical data. Do not enable .Mac for your administrator or system administrator (root) accounts.

You should not enable any options in the Sync pane of .Mac preferences.

You should not enable iDisk Syncing. If you must use a Public Folder, enable password protection.

You should not register any computers for synchronization in the Advanced pane of .Mac preferences.

To securely configure .Mac preferences:

1. Open .Mac preferences.
2. Deselect "Synchronize with .Mac."
3. Don't enable iDisk Syncing in the iDisk pane.
4. Don't register your computer for synchronization in the Advanced pane.

Securing Accounts Preferences

You can use Accounts preferences to perform two major security-related tasks: change or reset account passwords, and modify login options.

You should immediately change the password of the first account that was created on your computer. If you are an administrator, you can reset other user account passwords by selecting the account and clicking Change Password.

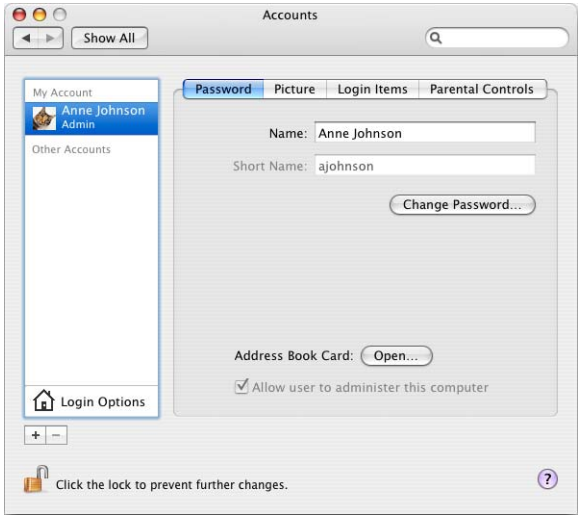

If your user account is managed by Open Directory, you should use Workgroup Manager to configure mobile accounts. For more information about mobile accounts and how to configure them, see "Managing Mobility Preferences" on page 161.

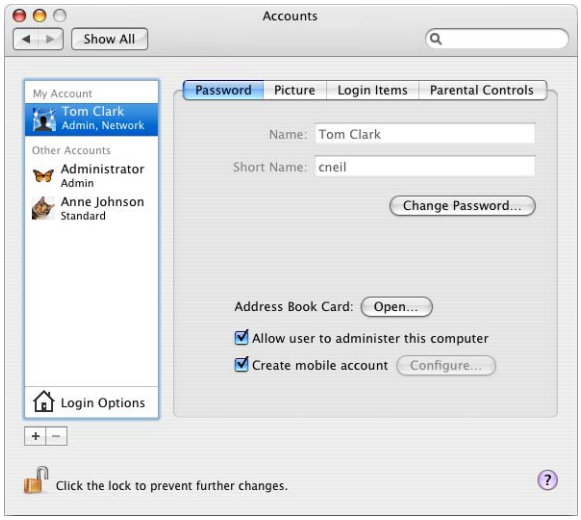

The password change and reset dialogs provide access to Password Assistant, an application that can analyze the strength of your chosen password and assist you in creating a more secure password. For information, see "Using Password Assistant" on page 73.

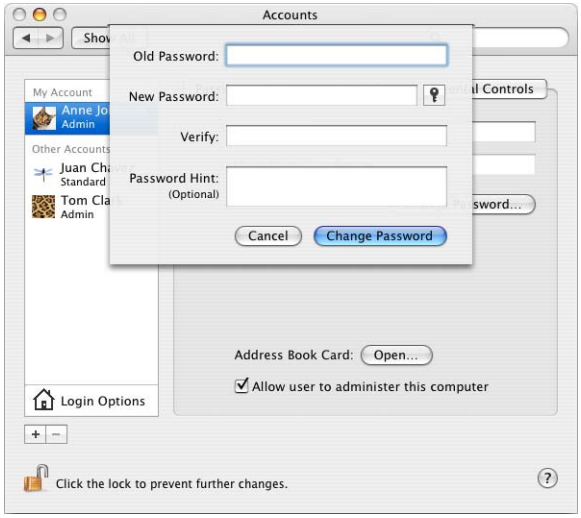

You should modify login options, so that you provide as little information as possible to the person trying to log on. You should require that the user know which account they want to log in with, and the password for that account. You shouldn't automatically log the user in, you should require that the user enter both a name and password, and that the user authenticate without the use of a password hint. Don't enable fast user switching—it is a security risk because it allows multiple users to be simultaneously logged in to the computer.

You should also modify login options to disable the Restart, Sleep, and Shut Down buttons. By disabling these buttons, the user cannot restart the computer without pressing the power key or logging in.

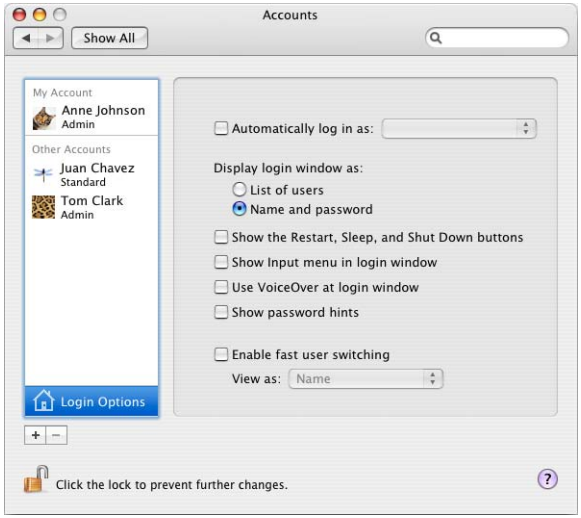

To securely configure Accounts preferences:

1 Open Accounts preferences.

2 Select an account and click the Password pane. Then, change the password by clicking Change Password.

A menu will appear asking you to input the old password, new password, verification of the new password, and a password hint. Do not enter a password hint, then click Change Password.

3 Click Login Options and select only Display login window as to name and password. Deselect all other options.

Securing Appearance Preferences

Recent items refer to applications, documents, and servers that you've recently used. You can access recent items by choosing Apple > Recent Items.

You should consider changing the number of recent items displayed in the Apple menu to none. If an intruder gains access to your computer, they can use recent items to quickly view your most recently accessed files. Additionally, intruders can use recent items to access any authentication mechanism for servers if the corresponding keychains are unlocked. Removing recent items provides a minimal increase in security, but it can deter very unsophisticated intruders.

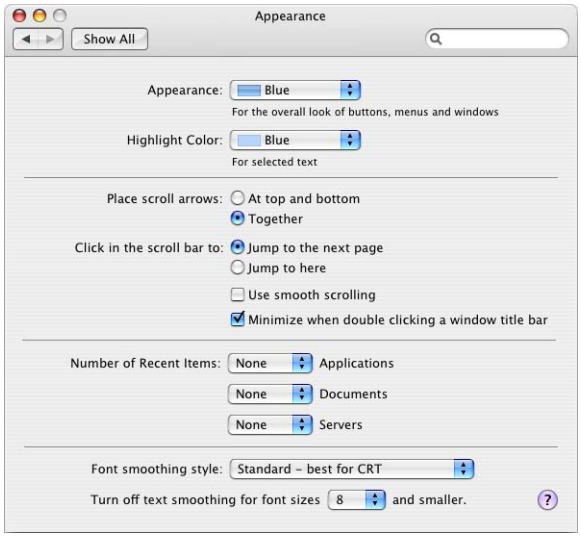

To securely configure Appearance preferences:

1. Open Appearance preferences.
2. Set all of the "Number of Recent Items" preferences to none.

Securing Bluetooth Preferences

Bluetooth allows wireless devices, such as keyboards, mice, and mobile phones, to communicate with the computer. If the computer has Bluetooth capability, Bluetooth preferences become available. If you don't see Bluetooth preferences, you cannot use Bluetooth.

Note: Some high-security areas do not allow RF communication. Consult your organization's requirements regarding Bluetooth settings.

When you disable Bluetooth in System Preferences, you must disable Bluetooth for every user account on the computer. This does not prevent users from reenabling Bluetooth. It is possible to restrict a user account's privileges so that the user cannot reenable Bluetooth, but to do this, you also remove several important user abilities, like the user's ability to change his or her own password. For more information, see "Securing Local Nonadministrator Accounts" on page 67.

To securely configure Bluetooth preferences:

1. Open Bluetooth preferences.
2. Set Bluetooth Power to Off.

Securing CDs & DVDs Preferences

The computer should not perform automatic actions when inserting CDs or DVDs. When you disable automatic actions in System Preferences, you must disable these actions for every user account on the computer. This does not prevent users from reenabling automatic actions. To prevent the user from reenabling automatic actions, you must restrict the user's account, so that the user cannot open System Preferences. For more information about restricting accounts, see "Securing Local Nonadministrator Accounts" on page 67.

To securely configure CDs & DVDs preferences:

1. Open CDs & DVDs preferences.
2. Choose Ignore for each pop-up menu to disable automatic actions when inserting media.

Securing Classic Preferences

Mac OS X includes an adaptation of Mac OS 9, known as "Classic." Mac OS 9 should be removed from the computer. If you remove Mac OS 9 and do not plan on using it, you do not need to configure Classic preferences. For instructions on how to remove Mac OS 9, see "Removing Mac OS 9" on page 51.

If you are going to use Mac OS 9 from a CD, DVD, or disk image, you must configure Classic preferences. Although Mac OS 9 has security issues that you cannot prevent, you can minimize Mac OS 9's security risks. For more information, see "Running Mac OS 9 from a Disc Image" on page 53.

In the Start/Stop pane of Classic preferences, do not set Classic to start when you log in, and do not set Classic to hide while starting. Mac OS X should also warn before starting Classic, and show Classic status in the menu bar. By changing these settings, you increase awareness when running Classic.

Turn off extensions in the Advanced pane of Classic preferences. Although Classic is not allowed to interact directly with hardware, you might have several extensions that are related to hardware and are therefore unnecessary.

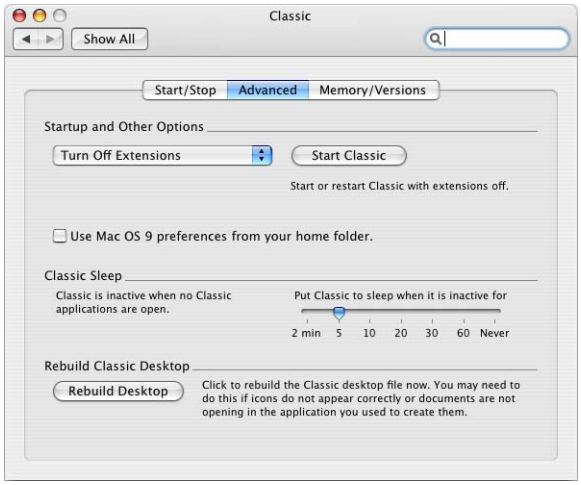

You can also use the Memory/Versions pane of Classic preferences to view the applications running in Mac OS 9. By choosing to show background applications, you become more aware of any malicious applications running in Mac OS 9.

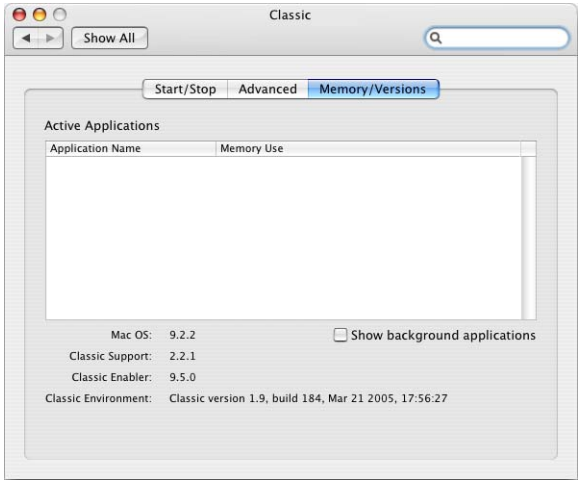

To securely configure Classic preferences:

1. Open Classic preferences.
2. In the Start/Stop pane, deselect "Start Classic when you login" and "Hide Classic while starting."
3. Select "Warn before starting Classic."
4. Click the Advanced pane, and select "Turn Off Extensions."

Securing Dashboard and Exposé Preferences

Your computer should require authentication when waking from sleep or screen saver. You can configure Dashboard & Exposé preferences to allow you to quickly start the screen saver if you move your mouse cursor to a corner of the screen. You should not configure any corner to disable the screen saver.

For information about requiring authentication for the screen saver, see "Securing Security Preferences" on page 104.

The Dashboard widgets included with Mac OS X Server can be trusted. However, you should be careful when you install third-party Dashboard widgets. You can install Dashboard widgets without having to authenticate. If you want to prevent Dashboard from running, set the keyboard and mouse shortcuts to "–."

When you configure Dashboard and Exposé preferences, you must configure these preferences for every user account on the computer. This does not prevent users from reconfiguring their preferences. It is possible to restrict a user account's privileges so that the user cannot reconfigure preferences. To do this, you will also remove several important user abilities, like the user's ability to change his or her own password. For more information, see "Types of User Accounts" on page 63.

If your organization does not want to use Dashboard because of its potential security risk, you can disable it from opening when selected.This is done from the command line, either locally or remotely.

To disable Dashboard locally from the command line:

1 Open Terminal.

2 Enter the command:

```
$ defaults write com.apple.dashboard mcx-disabled -boolean YES
```

This prevents Dashboard from opening.

3 Quit Terminal.

Securing Date & Time Preferences

Correct date and time settings are required for some authentication protocols, like Kerberos. Incorrect date and time settings can cause security issues.

The Date & Time preferences can automatically set the date and time based on a Network Time Protocol (NTP) server. If you require automatic date and time, use a trusted, internal NTP server.

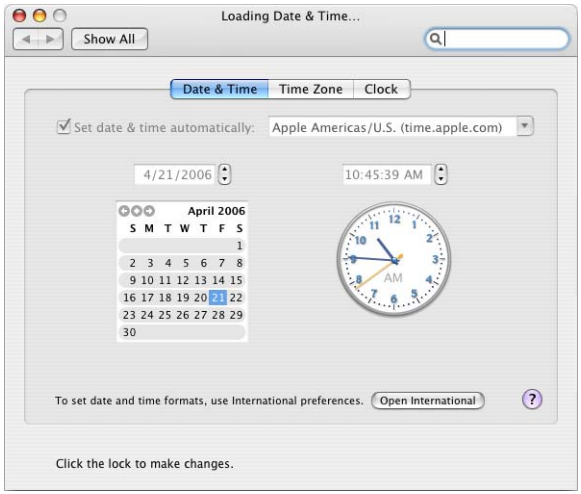

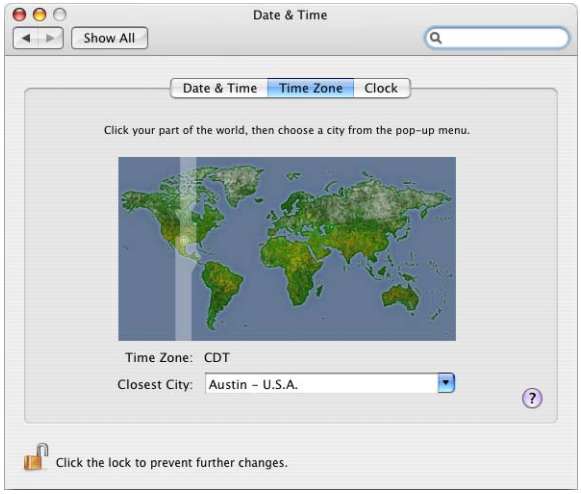

To securely configure Date & Time preferences:

1 Open Date & Time preferences.

2 In the Date & Time pane, enter a secure and trusted NTP server in the "Set date & time automatically" field.

3 In the Time Zone pane, choose a time zone.

Securing Desktop & Screen Saver Preferences

You can configure a password-protected screen saver to help prevent accessing of unattended computers by unauthorized users. Different authentication methods can be used to unlock screen savers, which include digital tokens, smart cards, or biometric readers. You should set a short inactivity interval to decrease the amount of time the unattended computer spends unlocked.

For information about requiring authentication for screen savers, see “Securing Security Preferences” on page 104.

You can configure Desktop & Screen Saver preferences to allow you to quickly enable or disable screen savers if you move your mouse cursor to a corner of the screen. You should not configure any corner to disable screen savers. You can also do this by configuring Dashboard & Exposé preferences.

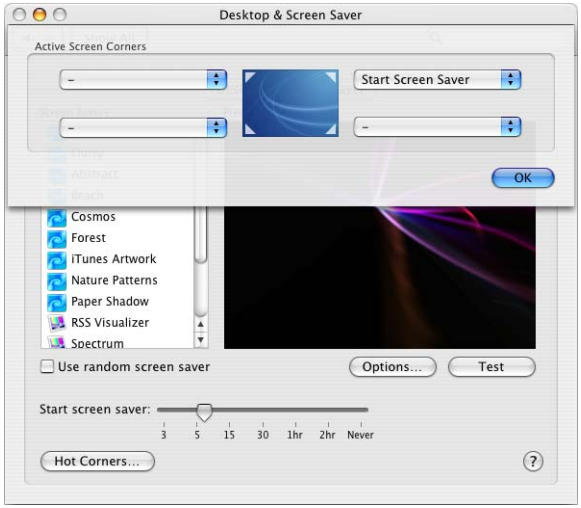

When you configure Desktop & Screen Saver preferences, you must configure these preferences for every user account on the computer. This doesn't prevent users from reconfiguring their preferences. It is possible to restrict a user account's privileges so that the user cannot reconfigure preferences. Doing this removes several important user abilities, like the user's ability to change his or her own password. For more information, see "Types of User Accounts" on page 63.

To securely configure Desktop & Screen Saver preferences:

1. Open Desktop & Screen Saver preferences.
2. Click the Screen Saver pane.
3. Set "Start screen saver" to a short inactivity time.
4. Click Hot Corners.
5. Set a corner to Start Screen Saver for quick enabling of the screen saver.

   Don't set any screen corner to Disable Screen Saver.

Securing Displays Preferences

If you have multiple displays attached to your system, be aware that enabling display mirroring might inadvertently expose private data to others. Having this additional display provides extra opportunity for other to see private data.

Securing Dock Preferences

You can configure the Dock to be hidden when not in use, which can prevent others from seeing what applications you have available on your computer when they pass by.

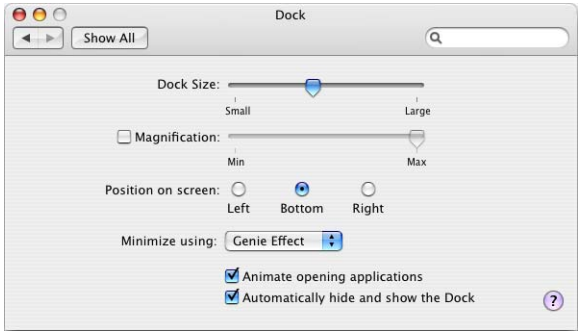

To securely configure Dock preferences:

1. Open Dock preferences.
2. Select "Automatically hide and show the Dock."

Securing Energy Saver Preferences

You can configure the period of inactivity required before a computer, display, or hard disk enters sleep mode, and require authentication by use of a password, digital token, smart card, or biometric reader when a user tries to use the computer. This is similar to using a password-protected screen saver. Mac OS X Server also allows you to set up different settings, depending on your power supply (power adapter or battery). For information about how to set up password protection for sleep mode, see "Securing Security Preferences" on page 104.

If you want to allow management and network visibility, you can configure the display and the hard disk to sleep, but not the computer.

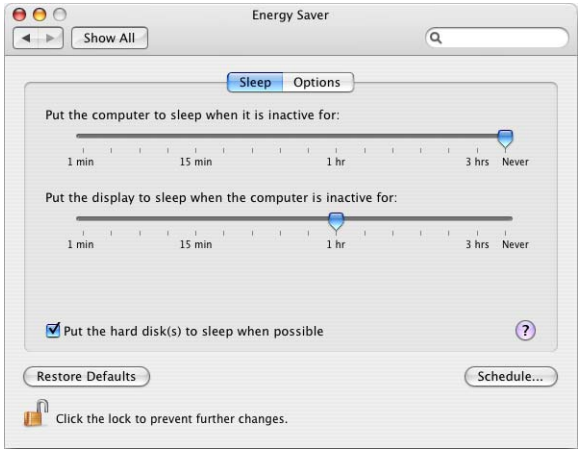

Also, the computer should not be set to restart after a power failure.

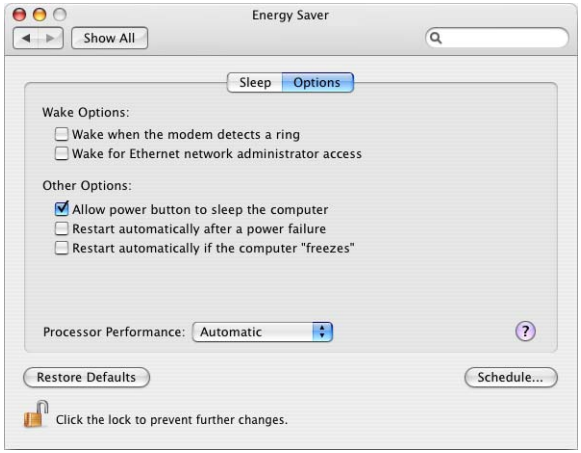

To securely configure Energy Saver preferences:

1. Open Energy Saver preferences.
2. Click the Sleep pane.
3. Set "Put the computer to sleep when it is inactive for" to Never.
4. Select "Put the hard drive disk(s) to sleep when possible."
5. Click the Options pane, and deselect "Wake when the modem detects a ring," "Wake from Ethernet network administrator access," "Restart automatically after a power failure," and "Restart automatically if the computer "freezes"."

Securing International Preferences

No security-related configuration is necessary. However, if your computer uses more than one language, check the security risk of the language character set. You should deselect any unused packages during the installation of Mac OS X Server.

Securing Keyboard & Mouse Preferences

Bluetooth should be turned off if not required. If Bluetooth is necessary, it is a good practice to disable allowing Bluetooth devices to wake this computer.

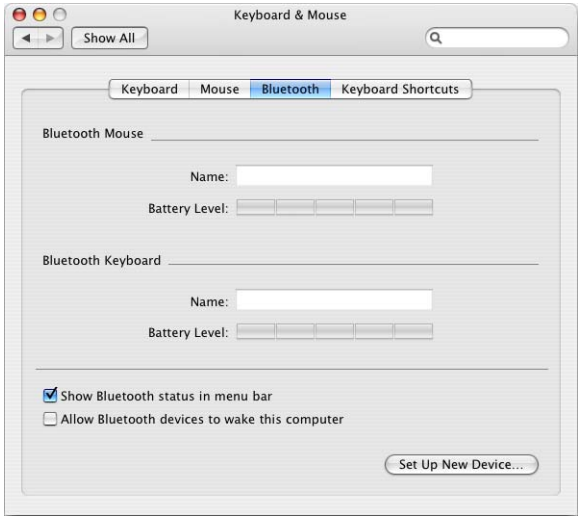

To securely configure Keyboard & Mouse preferences:

1. Open Keyboard & Mouse preferences.
2. Click Bluetooth.
3. Deselect "Allow Bluetooth devices to wake this computer."

Securing Network Preferences

You should disable any unused hardware devices listed in Network preferences. Enabled, unused devices (such as AirPort and Bluetooth) are a security risk.

Hardware is listed in Network preferences only if the hardware is installed in the computer.

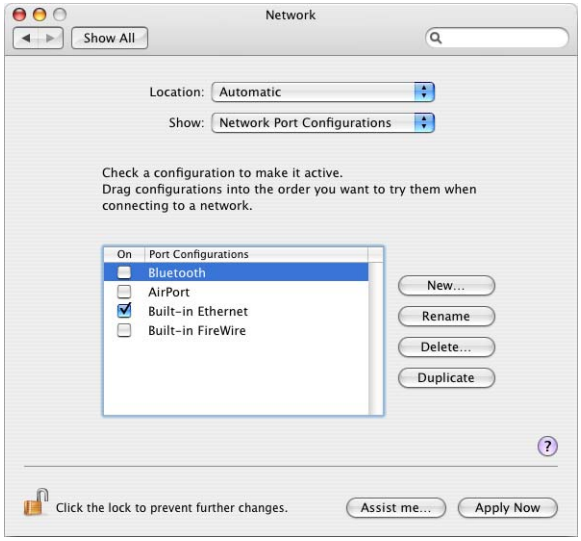

Some organizations use IPv6, a new version of the Internet Protocol (IP). The primary advantage of IPv6 is that it increases the address size from 32 bits (the current IPv4 standard) to 128 bits. An address size of 128 bits is large enough to support a huge number of addresses, even with the inefficiency of address assignment. This allows more addresses or nodes than are otherwise available. IPv6 also provides more ways to set up the address and simplifies auto configuration.

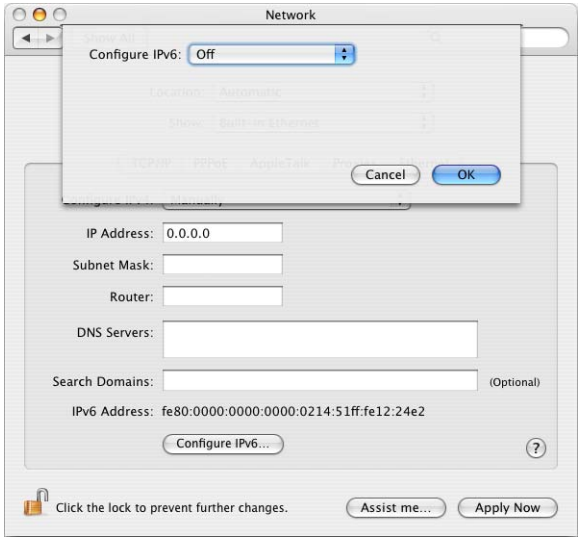

Unless IPv6 is specifically required by your organization, you should disable this capability by configuring IPv6 to be Off for each network port in use. By default, IPv6 is configured automatically.

To securely configure Network preferences:

1. Open Network preferences.
2. In the Show pop-up menu, choose your network device.
3. Click Configure IPv6.
4. In the Configure IPv6 pop-up menu, choose Off.
5. Click OK.
6. In the Show pop-up menu, choose Network Port Configurations.
7. Deselect any unused devices to disable them.

Securing Print & Fax Preferences

You should only use printers that are in a secure location. If you print confidential material in an insecure location, your confidential data sent to a printer might be viewable by unauthorized users. You should also be careful not to print to a shared printer, since that allows another computer to capture the complete print job directly. The remote computer could be maliciously monitoring and capturing confidential data being sent to the printer.

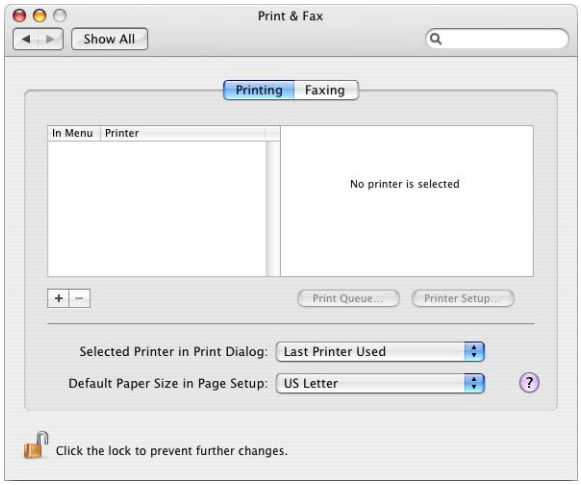

You should not send or receive faxes on your computer. By disabling faxes, you remove an additional avenue for potential attack.

To securely configure Print & Fax preferences:

1 Open Print & Fax preferences.

2 In the Faxing pane, deselect "Receive faxes on this computer."

Securing QuickTime Preferences

You should only download QuickTime movies from trusted, secure sources. By default, QuickTime stores downloaded movies in a cache. If someone gained access to your account, they would be able to see your previously viewed movies, even if you did not explicitly save them as files. You can change QuickTime preferences to disable the storing of movies in a cache.

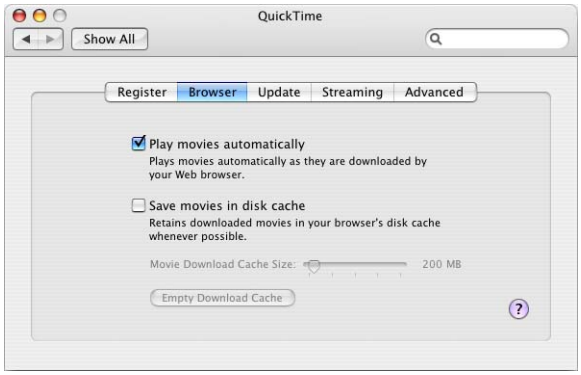

You should not install third-party QuickTime software unless you specifically require that software.

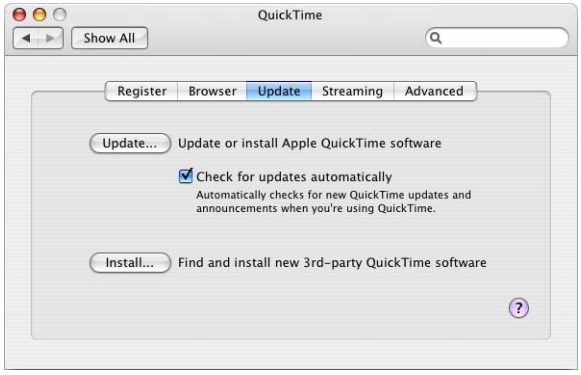

To securely configure QuickTime preferences:

1 Open QuickTime preferences.

2 In the Browser pane, deselect "Save movies in disk cache."

Securing Security Preferences

The settings in Security preferences cover a wide range of Mac OS X security issues.

Mac OS X includes FileVault, which encrypts the information in your home folder. FileVault uses the latest government-approved encryption standard, the Advanced Encryption Standard with 128-bit keys (AES-128). For more information about FileVault, see "Encrypting Home Folders" on page 120.

You should require a password to wake the computer from sleep or screen saver. This helps prevent unauthorized access to unattended computers. Although there is a lock button for Security preferences, individual users don't need to be authorized as an administrator to change this setting. You should enable this setting for every user on the computer.

To securely configure Security preferences:

1 Open Security preferences.

2 Select "Require password to wake this computer from sleep or screen saver."

3 Click "Turn On FileVault."

4 Authenticate with your account password.

5 Select "Use secure erase."

6 Click "Turn On FileVault."

7 Restart the computer.

Securing Sharing Preferences

By default, Remote Login is enabled while Apple Remote Desktop and Remote Apple Events are disabled in Sharing preferences. You should not enable any of these services unless you are required to use them. The following services are described in greater detail in Chapter 10, "Securing Remote Access Services," on page 191.

Service	Description
Remote Login	Allows users to access the computer remotely by using SSH. If you require the ability to perform remote login, SSH is more secure than telnet, which is enabled by default.
Apple Remote Desktop	Allows the computer to be accessed using Apple Remote Desktop.
Remote Apple Events	Allows the computer to receive Apple-specific events from other computers.

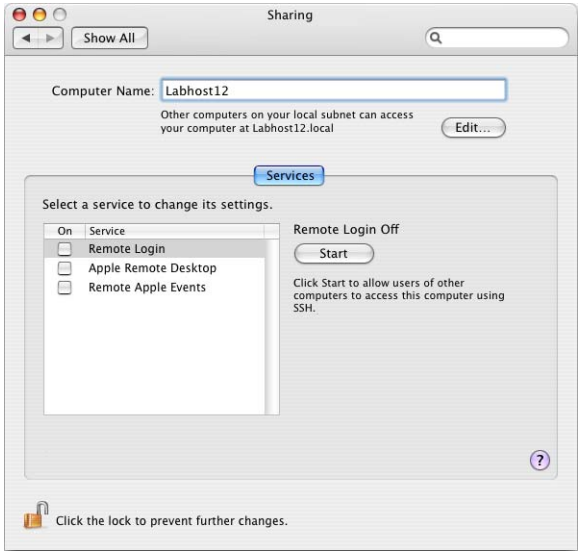

You can change your computer's name in Sharing preferences. When other users use Bonjour to discover your available services, your computer is displayed as *hostname*.local. To increase your privacy, change your computer's name so it does not indicate the purpose of the computer. Don't use the word "server" as the name or part of the name.

To securely configure Sharing preferences:

1. Open Sharing preferences.
2. Deselect Remote Login.
3. Change the default "Computer Name" to a name that does not indicate the purpose of the computer.

Securing Software Update Preferences

Your Software Update preferences configuration primarily depends on your organization's policy. For example, if your operational computer is connected to a managed network, the management settings determine what software update server to use.

Instead of using Software Update, you can also manually update your computer by using installer packages. You could install and verify updates on a test-bed computer before installing them on your operational computer. For more information about how to verify the authenticity of the installer packages and manually update your computer, see "Updating Manually from Installer Packages" on page 44.

After transferring installer packages to your computer, you should verify the authenticity of the installer packages. For more information, see "Verifying the Integrity of Software" on page 45.

When you try to install a software update, either by using Software Update or by using an installer package, you are required to authenticate with an administrator's name and password. This reduces the chance of accidental or malicious installation of software updates. Software Update will not install a software package that has not been digitally signed by Apple.

To securely configure Software Updates preferences:

1. Open Software Update preferences.
2. Click the Update Software pane.
3. Deselect "Check for updates" and "Download important updates in the background."

Securing Sound Preferences

Many Apple computers include an internal microphone, which can cause security issues. You can use Sound preferences to disable the internal microphone and the audio line-in port.

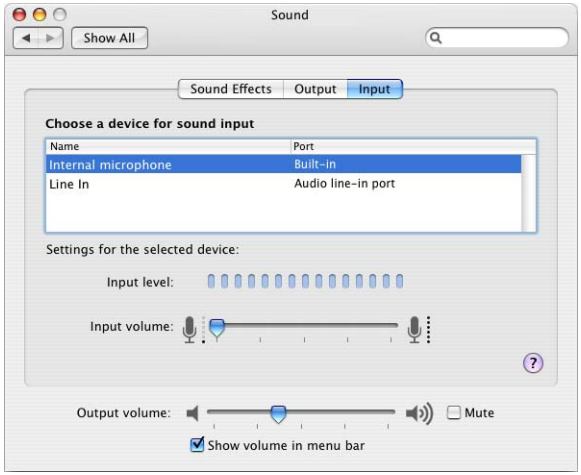

To securely configure Sound preferences:

1. Open Sound preferences.
2. Select Internal microphone (if present), and set "Input volume" to zero.
3. Select Line In, and set "Input volume" to zero.

   This ensures that "Line In" is the device selected rather than the internal microphone when preferences is closed, providing protection against inadvertent use of the internal microphone.

Securing Speech Preferences

Mac OS X includes speech recognition and text to speech features. You should only enable these features if you're working in a secure environment where no one else can hear you speak to the computer, or hear the computer speak to you. Also make sure that there are no audio recording devices that can record your communication with the computer.

If you do enable the text to speech feature, use headphones to keep others from overhearing your computer.

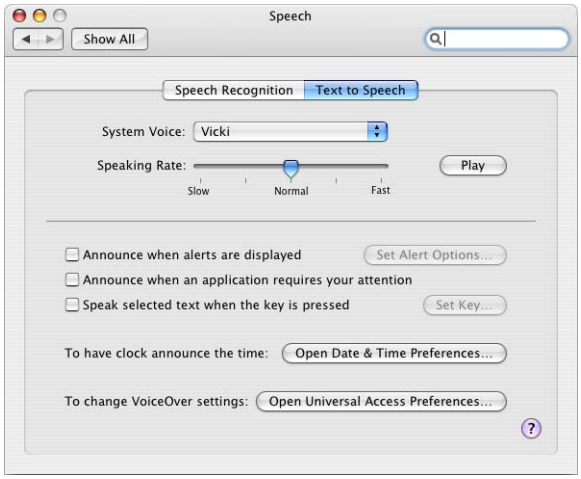

To securely configure Speech preferences:

1. Open Speech preferences.
2. Click the Speech Recognition pane, and set Speakable Items on or off.

   Change the setting according to your environment.
3. Click the Text to Speech pane, and change the settings according to your environment.

Securing Spotlight Preferences

Spotlight is a new feature in Mac OS X version 10.4. You can use Spotlight to search your entire computer for files. Spotlight searches not only the name and meta-information associated with each file, but also the contents of each file. Spotlight nullifies the use of file placement as an additional layer of security. You must still properly set access permissions on folders containing confidential files. For more information about access permissions, see "Repairing Disk Permission" on page 45.

By default, all categories are available for searching using Spotlight.

By placing specific folders or disks in the Privacy pane, you can prevent Spotlight from searching them. You should disable searching of all folders that contain confidential information. Consider disabling top-level folders. For example, if you store confidential documents in subfolders of ~/Documents/, instead of disabling each individual folder, disable ~/Documents/.

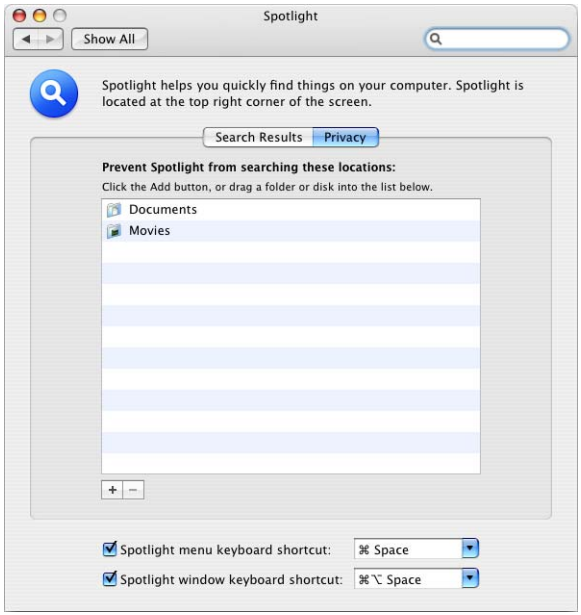

To securely configure Spotlight preferences:

1. Open Spotlight preferences.
2. In the Search Results pane, deselect any categories you don't want searchable by spotlight.
3. Click the Privacy pane.
4. Click the Add (+) button, or drag a folder or disk into the Privacy pane.
5. Folders and disks in the Privacy pane are not searchable by Spotlight.

You can use the `mdutil` tool turn Spotlight indexing off for a volume. For example, to erase the current meta data store and turn indexing off for a volume called *`volumename`*:

```
$ mdutil -E -i off volumename
```

For information, see the `mdutil` man page.

Securing Startup Disk Preferences

You can use Startup Disk preferences to make your computer start up from a CD, a network volume, a different disk or disk partition, or another operating system.

Be careful when selecting a startup volume. Choosing a network install image reinstalls your operating system and might erase the contents of your hard disk. If you choose a FireWire volume, your computer will start up from the FireWire drive plugged into the current FireWire port for that volume. If you connect a new, different FireWire drive to that FireWire port, your computer will startup from the first valid Mac OS X Server volume on the computer. This is assuming you have not enabled the Open Firmware password.

When you enable an Open Firmware password, the FireWire volume you selected is the only volume that will start the system. Open Firmware locks in the FireWire Bridge Chip GUID as a startup volume instead of the hard drive's GUID (as is done with internal hard drives). If the drive inside the FireWire drive enclosure is replaced with a new drive, the system can start from the new drive without having to bypass the Open Firmware password. To avoid this type of intrusion, make sure your hardware is physically secured. Open Firmware can also have a list of FireWire volumes that are approved for system startup.

For information about physically protecting your system, see "Protecting Hardware" on page 49.

You can also restart in target disk mode from Startup Disk preferences. When your computer is in target disk mode, another computer can connect your computer and access your computer's hard drive. The other computer has full access to all the files on your computer. All file permissions for your computer are disabled in target disk mode.

If you hold down the T key during startup, you enter target disk mode. You can prevent the startup shortcut for target disk mode by enabling an Open Firmware or EFI password. If you enable an Open Firmware or EFI password, you can still restart in target disk mode using Startup Disk preferences. For more information about enabling an Open Firmware or EFI password, see "Configuring Open Firmware Settings" on page 56.

To select a Startup Disk:

1. Open Startup Disk preferences.
2. Select a volume to use to start up your computer.
3. Click "Restart" to restart from the selected volume.

Securing Universal Access Preferences

Universal Access preferences are disabled by default. If you don't use an assistive device, there are no security-related issues. However, if you do use an assistive device, follow these guidelines:

- See the device manual for prevention of possible security risks.
- Enabling "VoiceOver" configures the system to read the contents under the cursor out loud, which might inadvertently disclose confidential data.
- These devices allow access to the computer that could reveal information in a compromising manner.

Securing Data and Using Encryption

6

Your data is the most valuable part of your computer. By using encryption, you can protect your data in the case of an attack or theft of your computer.

By setting global permissions, encrypting home folders, and encrypting portable data, you can be sure that your data is secure. By using the secure erase feature of Mac OS X Server, any deleted data is completely erased from the computer.

Understanding Permissions

Files and folders are protected by setting permission that restrict or allow users access to them. Mac OS X Server supports two methods of setting file and folder permissions:

- Portable Operating System Interface (POSIX) permissions—standard for UNIX operating systems.
- Access Control Lists (ACLs) permissions—used by Mac OS X, and compatible with Microsoft Windows Server 2003 and Microsoft Windows XP.

ACL uses POSIX in its process of verifying file and folder permissions. The process that ACL uses to determine if an action is allowed or denied, includes checking specific rules called access control entries (ACEs). If none of the ACEs apply, then standard POSIX permissions are used to determine access.

Note: In this guide, the term "privileges" refers to the combination of ownership and permissions, while the term "permissions" refers only to the permission settings that each user category can have (Read & Write, Read Only, Write Only, and None).

Setting POSIX Permissions

Mac OS X Server bases file permissions on POSIX standard permissions, such as file ownership and access. Every share point, file, or folder has read, write, and execute permission defined for three different categories of users (owner, group, and everyone). There are four types of standard POSIX access permissions that you can assign to a share point, folder, or file: Read & Write, Read Only, Write Only, and None.

Viewing POSIX Permissions

You can assign standard POSIX access permissions to these three categories of users:

- Owner—A user who creates a new item (file or folder) on the server is its owner and automatically has Read & Write permissions for that folder. By default, the owner of an item and the server administrator are the only users who can change its access privileges (allow a group or everyone to use the item). The administrator can also transfer ownership of the shared item to another user.
- Group—You can put users who need the same access to files and folders into group accounts. Only one group can be assigned access permissions to a shared item. For more information about creating groups, see the user management guide.
- Everyone—Any user who can log in to the file server: registered users and guests.

Before setting or changing POSIX permissions, you should view the current permission settings.

To view the permission of folders or files:

1. Open Terminal.
2. Run the `ls` command in Terminal.

```
$ ls -l
```

Output similar to the following will appear:

```
computer:~/Documents ajohnson$ ls -l
total 500
drwxr-xr-x 2 ajohnson ajohnson    68 Apr 28 2006 NewFolder
-rw-r--r-- 1 ajohnson ajohnson 43008 Apr 14 2006 file.txt
```

Note: The "'~" refers to your home folder, which in this case is /Users/ajohnson/. ~/Documents/ is the current working folder.

You can also use the Finder to view POSIX permissions. In the Finder, select the file, and choose File > Get Info. Open the Ownership & Permissions disclosure triangle to view POSIX permissions.

Interpreting POSIX Permissions

POSIX permissions can be interpreted by reading the first ten bits of the long format output listed for a file or folder.

```
drwxr-xr-x 2 ajohnson ajohnson    68 Apr 28 2006 NewFolder
-rw-r--r-- 1 ajohnson ajohnson 43008 Apr 14 2006 file.txt
```

In this example, the `NewFolder` has the POSIX permissions of `drwxr-xr-x` and has an owner and group of `ajohnson`. The `d` of the POSIX permissions signifies that `newfolder` is a folder. The first three letters after the `d (rwx)` signify that the owner has read, write, and execute permission for that folder. The next three characters, `r-x`, signify that group has read and execute permission. The last three characters, `r-x`, signify that all others have read and execute permission. In this example, any users who can access ajohnson's ~/Documents/ folder can also open the NewFolder folder and can view, but not modify or open, the file.txt file. "Read" POSIX permissions are propagated through the folder hierarchy. Although `NewFolder` has `drwxr-xr-x` privileges, only `ajohnson` will be able to access the folder. This is because ajohnson's ~/Documents/ folder has `drwx------` POSIX permissions.

By default, most of the user's folders have `drwx------` POSIX permissions. Only the ~/Sites/ and ~/Public/ folders have `drwxr-xr-x` permissions. This set of permissions allows other people to view folder contents without authenticating. You can change these folder permissions to `drwx------` if you do not want other people to view their contents. Within the ~/Public/ folder, the Drop Box folder has `drwx-wx-wx` POSIX permission. This allows users other than ajohnson to add files into a ajohnson's drop box but they are not able to view those files.

Occasionally, you'll see a `t` instead of an `x` for others' privileges on a folder used for collaboration. This `t` is sometimes known as the "sticky bit." Enabling the sticky bit on a folder prevents people from overwriting, renaming, or otherwise modifying other people's files. This is something that can become common if several people are granted `rwx` access. The sticky bit being set can appear as `t` or `T` depending on if the execute bit is set for others.

- If the execute bit appears as `t`, the sticky bit is set and has searchable and executable permissions.
- If the execute bit appears as `T`, the sticky bit is set, but does not have searchable or executable permissions.

For more information, see the `sticky` man page.

Modifying POSIX Permissions

After you determine the current POSIX permission settings, you can modify them using the `chmod` command.

To modify POSIX permissions:

1. Open Terminal.
2. Enter the following.

   ```
   $ chmod g+w file.txt
   ```

   This adds write permission for the group to file.txt.
3. View the permissions using the `ls` command.

   ```
   $ ls -l
   ```
4. Make sure the permissions are correct.

   ```
   computer:~/Documents ajohnson$ ls -l
   total 3
   drwxr-xr-x 2 ajohnson ajohnson    68 Apr 28 2006 NewFolder
   -rw-rw-r-- 1 ajohnson ajohnson 43008 Apr 14 2006 file.txt
   ```

For more information, see the `chmod` man page.

Setting File and Folder Flags

Files and folders can also be protected using flags. These flags, or permission extensions override standard POSIX permissions. These can be used to prevent the system administrator (root) from modifying or deleting files or folders.

Use the `chflags` command to enable and disable flags. The flag can only be set or unset by the file's owner or an administrator using `sudo`.

Viewing Flags

Before setting or changing file or folder flags, you should view the current flag settings.

To display flags set on a folder:

```
$ ls -lo secret
-rw-r--r-- 1 ajohnson ajohnson uchg 0 Mar  1 07:54 secret
```

In this example, the flag settings for a folder named *secret* are displayed.

Modifying Flags

After you determine the current file or folder flag settings, you can modify them using the `chflags` command.

To lock a folder using flags:

```
$ sudo chflags uchg secret
```

In this example, the folder named *secret* is locked. To unlock the folder, change `uchg` to `nouchg`.

> ***WARNING:*** There is an `schg` option available for the `chflags` command that sets the system immutable flag. This setting can only be undone when the computer is in single user mode.

For more information, see the `chflags` man page.

Setting ACL Permissions

For greater flexibility in configuring and managing file permissions, Mac OS X Server implements access control lists (ACL). An ACL is an ordered list of rules that control file permissions. Each rule or access control entry (ACE) contains the following components:

- User—owner, group, and other
- Action—read, write, or execute
- Permission—allow or deny the action

The rules specify the permissions to be granted or denied to a group or user, and how these permissions are propagated throughout a folder hierarchy.

Setting ACL Permissions Using Workgroup Manager

ACLs in Mac OS X Server let you set file and folder access permissions for multiple users and groups, in addition to the standard POSIX permissions. This makes it easy to set up collaborative environments with smooth file sharing and uninterrupted workflows, without compromising security. Mac OS X Server has implemented file system ACLs that are fully compatible with Microsoft Windows Server 2003 and Windows XP.

To determine if an action is allowed or denied, the ACEs are considered in order. The first ACE that applies to a user and action determines the permission and no further ACEs are evaluated. If none of the ACEs apply, then standard POSIX permissions determine access.

To set ACL permissions using Workgroup Manager:

1. Open Workgroup Manager and click Sharing.
2. Make sure ACLs are enabled for the volume on which the share point or folder is located.
3. To enable ACLs for a volume, click All, select the volume, select "Enable Access Control Lists on this volume" in the General pane, and click OK when prompted.

4 Click Save.

5 Click All and select the share point or folder.

6 Click Access.

7 Click Users & Groups to open the Users & Groups drawer.

8 Drag groups and users in the order you want them in the Access Control List.

9 To edit the ACEs, select the entry from the Access Control List.

10 Click the "Edit selected item" (/) button.

11 Select the permission type from the pop-up menu.

12 Select desired permissions in the Permissions list.

13 Click OK.

14 Click Save.

Setting ACL Permissions for a File

You can set ACL permissions for files. The `chmod` command enables an administrator to grant specific users reading, writing, and editing privileges for a single file.

To set ACL permissions for a file:

1 Allow required users access to specific files.

For example, to allow Anne Johnson access to read a specific file secret.txt, enter the following in Terminal:

```
$ chmod +a "ajohnson allow read" secret.txt
```

2 Allow required groups of users access to specific files.

For example, to allow the engineers group access to delete the file secret.txt, enter the following in Terminal:

```
$ chmod +a "engineers allow delete" secret.txt
```

3 Deny access privileges to specific files.

For example, to prevent Tom Clark from modifying the file secret.txt, enter the following in Terminal.

```
$ chmod +a "tclark deny write" secret.txt
```

4 View and validate the ACL modifications with the `ls` command.

```
$ ls -le secret.txt
-rw------- 1 ajohnson admin 43008 Apr 14 2006 secret.txt
0: ajohnson allow read
1: tclark deny write
2: engineers allow delete
```

For more information, see the `chmod` man page.

Setting Global File Permissions

Every file or folder has POSIX permissions associated with it. When you create a new file or folder, the umask setting determines these POSIX permissions. The default umask setting 022 (in hexidecimal), removes group and other write permissions. Group members and other users can read and run these files or folders. If you change this umask setting to 027, files and folders can still be read and run by group members, but cannot be accessed in any way by others. If you want to be the only user who can access your files and folders, set the umask setting to 077.

To change the globally defined umask setting, change the NSUmask setting. However, not all applications recognize the NSUmask setting. Therefore, there is no guarantee that files and folders created by other applications will have proper umask settings. The NSUmask setting also doesn't affect some command-line tools.

Users can use the Finder's Get Info window or the `chmod` command-line tool to change permissions for individual files and folders.

> ***WARNING:*** Many installations depend on the default umask settings. There can be unintended and possibly severe consequences to changing it. Instead, use inherited permissions, which are applied simply by setting permissions on a folder. All files contained within that folder will inherit the permissions of that folder.

To set the global umask:

1. Open Terminal.
2. Change the NSUmask setting to be the decimal equivalent of the umask setting:

   ```
   $ sudo defaults write /Library/Preferences/.GlobalPreferences NSUmask 23
   ```

   You must be logged in as a user who can use `sudo` to perform these operations. This example sets the global umask to 027, which has the decimal equivalent of 23. Replace `23` with the decimal equivalent of your desired umask setting. This command requires that you use the decimal equivalent, and not a hexidecimal number.

   Important: Make sure the path you enter is .GlobalPreferences—not .GlobalPreferences.plist, which might be accidentally added by Terminal's auto completion feature.
3. Log out.

Changes to umask settings take effect at the next login. Users can use the Finder's Get Info window or the `chmod` command-line tool to change access settings for individual files and folders.

Securing Your Home Folder

Change the permissions of each user's home folder so that they are no longer world-readable or world-searchable. When FileVault is not enabled, the permissions on the home folder of a newly-created user account allow any other user to browse its contents. The ~/Public and ~/Public/Drop Box folders within each home folder require these permissions. However, users may inadvertently save sensitive files directly into their home folder, instead of into the more-protected ~/Documents, ~/Library, or ~/Desktop folders. Although ~/Public and ~/Public/Drop Box folders will no longer work as intended, the permissions on each user's home folder should be changed to prevent other users from browsing its contents.

Enter the following command to change home folder permissions:

```
$ sudo chmod 750 /Users/username
```

Replace *username* with the name of the account.

Run this command immediately after everytime someone creates a new account. The 750 permission setting still allows members of the group owning the folder to browse it, but in Mac OS X version 10.3 or later that group consists only of the user. If more advanced group management is performed and members of the group owning the folder should not be granted permission to browse it, then the command above should be issued with the permission 700 instead of 750. The user, as the owner of his home folder, can alter its permission settings at any time, and can change these settings back.

Encrypting Home Folders

Mac OS X Server includes FileVault, which can encrypt your home folder and all the files contained within it. You should use FileVault on portable computers, and on any other computers whose physical security you cannot guarantee. You should enable FileVault encryption for your computer and for all its user accounts.

FileVault moves all the content of your home folder into a sparse disk image that uses AES-128 encryption. The sparse format allows the image to maintain a size proportional to its contents, which can save disk space.

If you remove files from a FileVault-protected home folder, it takes some time to recover free space from the home folder. Once optimized, you can access files in FileVault-protected home folders without noticeable delays. If you're working with confidential files that you plan to later erase, store those files in separate encrypted images that are not located in your home folder. You can then erase those images without having to recover free space. For more information, see "Encrypting Portable Files" on page 123.

If you've insecurely deleted files before using FileVault, these files are still recoverable after activating it. When initially enabling FileVault, securely erase free space. For information, see "Securely Erasing Data" on page 125.

FileVault does not encrypt or protect files transferred over the network or saved to removable media, so you'll want to encrypt specific files or folders. If you mount these encrypted images, all data transmitted over the network will be encrypted with AES-128. For instructions about how to encrypt specific files or folders for transfer from your network home folder, see "Encrypting Portable Files" on page 123.

To set up FileVault, you create a master password. If you forget your login password, you can use your master password to recover encrypted data. If you forget both your login password and your master password, you cannot recover your data. Consider sealing your master password in an envelope and storing it in a secure location. You can also use Password Assistant to help create a complex master password that cannot be easily compromised. For information, see "Using Password Assistant" on page 73 and "Creating Complex Passwords" on page 298.

Enabling FileVault copies all data from your home folder into an encrypted home folder. After copying, FileVault erases the unencrypted data. By default, FileVault insecurely erases the unencrypted data, but you have the option of using secure erase. Enable secure erase, so that your unencrypted data is securely erased.

Using FileVault Master Keychain

A FileVault master keychain can be set to decrypt any account using FileVault to encrypt data. FileVault keychain should be set, to ensure data is not lost in the event of a forgotten password. If you forget the FileVault account password, which is used to decrypt their encrypted data, the FileVault master keychain is used to decrypt the data.

To create the FileVault master keychain:

1. Open Security preferences.
2. Click Master Password and set a master password.

   Select a very strong password and consider splitting the password into at least two components (first half/second half). Using Password Assistant can ensure that the quality of the password selected is strong. Each password component is kept by separate security administrators to avoid one person knowing the full password.

   This prevents a single person from unlocking (decrypting) a FileVault account by requiring two or more security administrators. For more information, see "Using Password Assistant" on page 73.

   This creates a keychain called FileVaultMaster.keychain located in /Library/Keychains/. The FileVault master keychain now contains both a FileVault recovery key (self-signed root CA certificate) and a FileVault master password key (private key).

3 You can delete the corresponding certificate called FileVaultMaster.cer, located in the same location as the FileVaultMaster.keychain.

   FileVaultMaster.cer is only used for importing the certificate into the keychain. This is only a certificate and does not contain the corresponding private key, so there are no security concerns with anyone gaining access to this certificate.

4 Make a copy of FileVaultMaster.keychain and put it in a secure place.

5 Delete the private key from FileVaultMaster.keychain created on the computer to modify the keychain.

   This ensures that even if someone is able to unlock the FileVault master keychain, they are unable to decrypt the contents of a FileVault account, since no FileVault master password private key is available for the decryption.

Centrally Managing FileVault

The modified FileVault master keychain can now be distributed to all of your network computers. This can be done by transferring FileVaultMaster.keychain to the desired computers using Apple Remote Desktop, a distributed installer executed on each computer, various scripting techniques, or just including it in the original disk image if your organization restores systems with a default image.

This provides network management of any FileVault account created on any computer with the modified FileVaultMaster.keychain located in the /Library/Keychains/ folder. These computers indicate that the master password is set in Security preferences.

When an account is created and the modified FileVault master keychain is present, the public key from the FileVault recovery key is used to encrypt the dynamically generated AES 128-bit symmetric key that is used for the encryption and decryption of the encrypted disk image (FileVault container).

To decrypt access to the encrypted disk image, the FileVault master password private key is required to decrypt the original dynamically generated AES 128-bit symmetric key. The user's original password continues to work as normal, but the assumption here is that the master password service is being used because the end user has forgotten the password or the organization must perform data recovery from a user's computer.

To recover a network managed FileVault system account:

1 Retrieve the copy of FileVaultMaster.keychain that was stored away before deleting the private key during modification.

2 Bring together all the security administrators involved in generating the master password. More than one individual is needed if the master password was split into two or more password components.

   Note: They must have root access to perform the restoring of the FileVaultMaster.keychain.

3 Restore the original keychain to the /Library/Keychains/ folder of the target computer replacing the installed one.

4 Ensure that the restored FileVaultMaster.keychain has the appropriate ownership and permissions set, similar to the following example.

```
-rw-r--r-- 1 root admin 24880 Mar 2 18:18 FileVaultMaster.keychain
```

5 Log in to the FileVault account you are attempting to recover and incorrectly enter the account password three times. If "Password Hints" is enabled, it then gives you an additional try after displaying the hint.

6 When prompted for the master password, the security administrators must combine their password components to unlock access to the account.

7 The account will be unlocked and you will be asked to provide a new password for the account which will also be used to encrypt the original symmetric key used to encrypt and decrypt the disk image.

Note: This process does not reencrypt the FileVault container, but simply reencrypts the original symmetric key with a key derived from the new master password you just entered.

8 You are now logged in to the account and given access to the user's home folder.

This process does not change the password used to protect the user's original login keychain, since that password is not known or stored anywhere. Instead, this process creates a new login keychain with the password just entered as the user's new account password.

Encrypting Portable Files

To protect files that you want to transfer over a network or save to removable media, you should either encrypt a disk image, or you should encrypt the individual files and folders. FileVault does not protect files transmitted over the network or saved to removable media.

Using a server-based encrypted disk image provides the added benefit of encrypting all network traffic between the computer and the server hosting the mounted encrypted disk image.

Creating a New Encrypted Disk Image

You can create a read/write or sparse image to encrypt and securely store data. A read/write image consumes the entire space that was defined when the image was created. For example, if the maximum size of a read/write image is set to 10 GB, then that image will consume 10 GB of space even if it contains only 2 GB of data. A sparse image will only consume the amount of space containing data in the image. For example, if the maximum size of a sparse image is 10 GB and the data contained in it is only 2 GB, it will consume only 2 GB of space.

If you are in a situation where it is possible to have unauthorized administrator access to your computer, creating an encrypted blank disk image is preferable to creating an encrypted disk image from existing data.

Any permissions set on an internal or external HFS+ hard disk that is created on one computer will be ignored by default, when mounted on another computer. This prevents possible conflicts with duplicated UID numbers. By using disk images to store data on an internal or external HFS+ hard disk, you can eliminate this permission vulnerability.

Creating an encrypting image from existing data copies the data from an unprotected area into the encrypted image. If the data is sensitive, it is better to create the image prior to creating the documents, since the working copies, backups, or caches of files would all be created in the encrypted storage from the start.

To create a new encrypted disk image:

1. Open Disk Utility.
2. Choose File > New > Blank Disk Image.
3. Enter and name for the image and choose where to store it.
4. Choose the size of the image by clicking the Size pop-up menu.

   You cannot increase the size of an image after creating it. Make sure that the size of the image is large enough for your needs.
5. Choose an encryption method by clicking the Encryption pop-up menu.

   AES-128 is a strong encryption algorithm.
6. Choose a format by clicking the Format pop-up menu.

   Although there is some overhead, the sparse format allows the image to maintain a size proportional to its contents (up to its maximum size), which can save disk space.
7. Click Create.
8. Enter a new password and verify it.

   You can easily access Password Assistant from this window. For more information, see "Using Password Assistant" on page 73.
9. Deselect "Remember password (add to Keychain)." Click OK.

Creating an Encrypted Disk Image from Existing Data

If you must maintain data confidentiality when transferring files from your computer, but you don't need to encrypt files on your computer, create a disk image from existing data. Such situations include unavoidable plain text file transfers across a network, such as email attachments or FTP, or copying to removable media, such as a CD-R or floppy disk.

If you plan to later add more files to this image, instead of creating an image from existing data, create a new encrypted disk image, and add your existing data to it. For more information, see "Creating a New Encrypted Disk Image" on page 123.

To create an encrypted disk image from existing data:

1. Open Disk Utility.
2. Choose File > New > Disk Image from Folder.
3. Select a folder, and click Image.
4. Choose File > New > Blank Disk Image.
5. Enter and name for the image and choose where to store it.
6. Choose a format by clicking the Format pop-up menu.

   The compressed disk image format helps you save hard disk space by reducing your disk image size.
7. Choose an encryption method by clicking the Encryption pop-up menu.

   AES-128 is a strong encryption algorithm.
8. Click Save.
9. Enter a new password and verify it.

   You can easily access Password Assistant from this window. For more information, see "Using Password Assistant" on page 73.
10. Deselect "Remember password (add to Keychain)." Click OK.

Securely Erasing Data

When you erase a file, you're actually just removing information that tells the file system where to find the file. The file's location on the disk is marked as free space. It is still possible to get this data from the disk if other files have not been written to the free space.

Mac OS X Server provides you with several ways to securely erase files:

- Zero-out erase—refers to setting all data bits on the disk to 0.
- 7-pass erase and 35-pass erase—use algorithms of varying complexity to overwrite the disk.

The zero-out erase is the quickest, while the 35-pass erase is the most secure, but it is also 35 times slower then the zero-out erase. Using secure erase ensures that no residual sensitive data remains on a drive or volume.

Each time you use a 7-pass or 35-pass secure erase, the following seven-step algorithm is used to prevent the data from ever being recovered:

- Overwrite file with a single character
- Overwrite file with zeroes
- Overwrite file with a single character
- Overwrite file with random characters
- Overwrite file with zeroes
- Overwrite file with a single character
- Overwrite file with random characters

Note: The 7-pass secure erase conforms to the Department of Defense standard for the sanitization of magnetic media. A 35-pass erase uses the extremely advanced Gutmann algorithm to help eliminate the possibility of data recovery.

Using Disk Utility to Securely Erase a Disk or Partition

You can use Disk Utility to securely erase an entire disk or free space from partitions.

> ***WARNING:*** Securely erasing a partition is irreversible. Be sure to back up any critical files that you want to keep before erasing the partition.

To securely erase a disk using Disk Utility:

1. Open Disk Utility (located in /Applications/Utilities/).
2. Select the partition or drive you want to securely erase.

   You cannot erase the start up partition only second partitions or secondary drives.
3. Click Security Options and select either the 7-pass erase or 35-pass erase to ensure that the data on the drive will be completely erased. Then click OK.
4. Choose either Mac OS Extended (Journaled) format or Mac OS Extended case-sensitive (Journaled) HFS+ format, which supports case-sensitive filenames (which is useful for legacy UNIX applications).
5. Choose "Erase" and the secure erase process begins.

   Secure erase can take a while to complete, depending on the amount of disk or partition space being erased and the method chosen.

Using Command-Line Tools to Securely Erase Files or Folders

You can use the `srm` command in Terminal to securely erase files or folders. By using `srm`, you have the flexibility to remove each specified file or folder by overwriting, renaming, and truncating the file or folder before erasing them. This prevents other people from undeleting or recovering any information about the file or folder.

For instance, `srm` supports simple methods, like overwriting data with a single pass of zeros, to more complex methods, like using a 7-pass erase or 35-pass erase. The `srm` command cannot remove write protected files owned by another user, regardless of the permissions on the folder containing the file.

> ***WARNING:*** Erasing files with `srm` is irreversible. Be sure to back up any critical files that you want to keep before securely erasing files.

To securely erase a folder named secret:

```
$ srm -r -s secret
```

The `-r` option removes the content of the folder, and the `-s` option (simple) only overwrites with a single random pass.

For a more secure erase, you can use the `-m` option (medium) to perform a 7-pass erase of the file. The `-s` option overrides the `-m` option, if both are present. If neither is specified, the 35-pass erase is used.

For more information, see the `srm` man page.

Using Secure Empty Trash

You can use the Secure Empty Trash menu option to quickly and securely erase all files stored in the Trash. This command uses a 7-pass erase. Depending on the total size of the files erased, securely emptying the trash might take some time to complete.

> ***WARNING:*** Using Secure Empty Trash is irreversible. Be sure to back up any critical files that you want to keep before securely erasing files.

To use Secure Empty Trash:

1. Open Finder.
2. Choose Finder > Secure Empty Trash.
3. Click OK.

Using Disk Utility to Securely Erase Free Space

You can use Disk Utility to securely erase free space on partitions, using a zero-out erase, a 7-pass erase, or a 35-pass erase.

To securely erase free space using Disk Utility:

1. In Finder, open Disk Utility located in the /Applications/Utilities/ folder.
2. Select the partition where you want to securely erase free space.

   Ensure that you select a partition, not a drive. Partitions are contained within drives and are indented one level in the list on the left.
3. Click "Erase" and then click "Erase Free Space."
4. Choose one of the erase options and click "Erase Free Space."

   Securely erasing free space can take a while to complete depending on the amount of free space being erased and the method chosen.
5. Choose Disk Utility > Quit Disk Utility.

Using Command-Line Tools to Securely Erase Free Space

You can securely erase free space from the command line by using the `diskutil` command. However, ownership of the affected disk is required. This tool allows you to securely erase using one of the three levels of secure erase:

- 1—zero-out secure erase (also known as single-pass)
- 2—7-pass secure erase
- 3—35-pass secure erase

To erase free space using a 7-pass secure erase (indicated by the number 2):

```
$ diskutil secureErase freespace 2 /dev/disk0s3
```

For more information about how to securely erase free space, see the `diskutil` man page.

Securing Accounts, Share Points, and Network Views

7

Use Workgroup Manager to set up and manage home folders, accounts, preferences, and settings for clients.

Mac OS X Server includes Workgroup Manager, a user management tool you can use to create and manage accounts, share points, and network views. When managing accounts, you can define core account settings like name, password, home folder location, and group membership. You can also manage preferences, allowing you to customize the user's experience, granting or restricting access to his or her own computer's settings and to network resources.

Workgroup Manager works closely with a directory domain. Directory domains are like databases, only specifically geared towards storing account information and handling authentication. For more information about Open Directory, see Chapter 17, "Securing Directory Services."

For information about using Workgroup Manager, see the user management guide.

Open Directory and Active Directory

Mac OS X Server supports both Open Directory and Active Directory domains for client authentication.

- Open Directory uses LDAPv3 as its connection protocol. LDAPv3 includes several security options, such as disabling clear-text passwords, encrypting all packets, and blocking man-in-the-middle attacks. For Windows clients, it is LAN manager NTLMv1 and NTLMv2, both of which are very weak password hashing schemes. For more information about how to configure these options, see "Configuring Open Directory Policies" on page 263.
- Active Directory connections are not as secure as Open Directory when all of its security settings are enabled. For example, users cannot receive directory services from an Active Directory server that enables digitally signing or encrypting all packets. Active Directory also allows the use of Highly Secure (HISEC) templates. Users can use third-party tools to further secure their Active Directory connections.

Users can mutually authenticate with both Open Directory and Active Directory. Both use Kerberos to authenticate. Kerberos is a ticket-based system that enables mutual authentication. The server must identify itself by providing a ticket to a users' computer. This prevents your computer from connecting to rogue servers. Users must enable trusted binding to mutually authenticate with Open Directory or Active Directory.

For more information about Open Directory and Active Directory, see the Open Directory administration guide.

Configuring Share Points

You can use the Sharing pane in Workgroup Manager to configure share points, which are hard disks (or hard disk partitions), CD-ROM discs, or folders that contain files you want users to share. You can use folders within these share points as home folder locations. You can set up share points so they can be accessed using file services. For more information, see Chapter 14, "Securing File Services," on page 235.

Using network home folders stored on a share point is inherently less secure than using local home folders. An intruder can access your network home folder through an insecure network connection.

NFS file access is not based on user authentication, but on the user ID and the client IP address, so it is generally less secure than AFP. Use NFS only if you need to provide home folders for a large number of users who use UNIX workstations.

To host home folders for Windows clients, use SMB/CIFS. SMB/CIFS is a protocol used by Windows to access share points. You can set up a share point for SMB/CIFS access only, so that Windows users have a network location for files that can't be used on other platforms. Like AFP, SMB/CIFS also requires authenticating with a valid user name and password to access files. SMB/CIFS uses NTLMv1 and NTLMv2 encryption, both of which are very weak password hashing schemes.

You cannot use FTP share points to host home folders. You should not use FTP share points. If you need to use FTP for file transfers, use the `sftp` command. The `sftp` command provides a secure means of authentication and data transfer while FTP does not. For more information about using the `sftp` command, see "Securing Remote Login" on page 191.

Check that all shares on local system drives are configured to grant access to only specific users or groups, and are not open to everyone. Removing any open shares will prevent unwanted access to your computer and prevent your computer being used to maliciously access additional computers on the network. Don't share files unnecessarily.

Configuring Workgroup Manager for Working with Share Points

You can enable Secure Socket Layer (SSL) transactions for working with share points in Workgroup Manager. You can also configure the refresh rate for the Sharing pane, which can help you test share point modifications.

To configure Workgroup Manager:

1. In Workgroup Manager, choose Workgroup Manager > Preferences.
2. Select "Use secure transactions (SSL) for Sharing."
3. If you want to increase the refresh rate for the Sharing pane, enter a value lower than the default value of 300 in the "Auto-refresh Sharing every # seconds" field.
4. Click OK.

Disabling Share Points

Disable any unused share points and also disable any unused sharing protocols. Enabled share points and sharing protocols can provide an additional avenue of attack for intruders.

If you disable all the share points using a particular sharing protocol, you should also disable that protocol. For more information, see "Configuring Windows File Sharing Service" on page 241.

To disable a share point:

1. In Workgroup Manager, click Sharing.
2. Click Share Points.
3. Select the share point that you want to disable.
4. Click General.
5. Deselect "Share this item and its contents."

Restricting Access to a Share Point

Before enabling a share point, restrict the access permissions for the folder that will act as the share point. Only allow users who must use the share point to be able to access it.

You can use Workgroup Manager's Sharing pane to set POSIX and ACL permissions to restrict share points to only being accessible by certain users. You can use a combination of the two permission types to provide a fine level of granularity for your users.

> *WARNING:* Carefully set access permissions. Incorrectly set access permissions can prevent legitimate users from accessing folders and files, or they can allow malicious users to access folders and files.

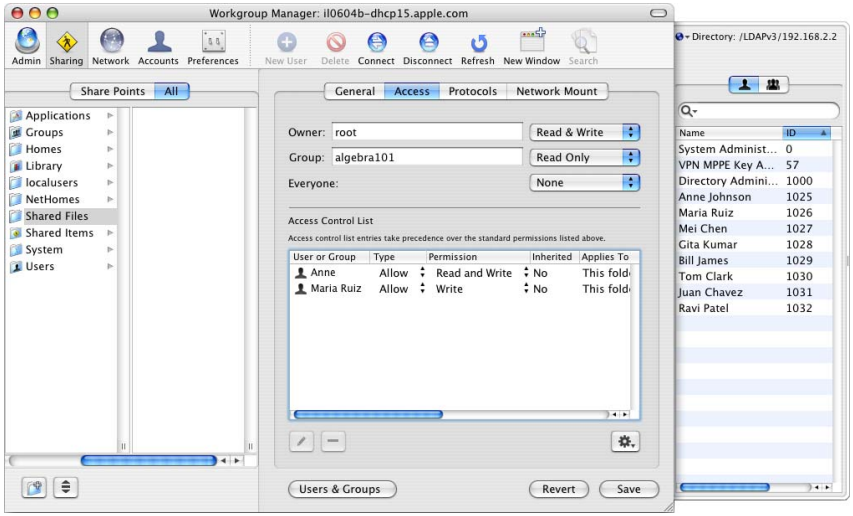

To restrict access to a share point:

1 In Workgroup Manager, click Sharing.

2 Click All.

3 Select the share point, or the folder that you want to use as a share point.

4 Click Access.

5 Click Users & Groups to display the Users & Groups drawer.

The Users & Groups drawer slides out from the right side of Workgroup Manager.

6 Click the small globe in the drawer and select your network directory domain.

7 To replace the POSIX owner and group, drag users or groups from the drawer into either the Owner or Group fields.

Make sure you understand the implications of changing a folder's owner and group. For more information, see "Setting POSIX Permissions" on page 113.

8 In the Owner, Group, and Everyone pop-up menus, choose the level of access that you want to grant.

If you're configuring a home folder's permissions, give the owner Read & Write privileges, but reduce group and everyone privileges to None.

The default for home folders is that the staff group and everyone have read privileges. All accounts are also members of the staff group. These two privileges allow everyone to view the contents of the home folder. If you want someone besides the owner being able to view the contents of the home folder, replace staff with that account.

9 To grant ACL permissions to a user or group, drag users or groups from the drawer into the Access Control List list.

By granting users ACL permissions, you override any applicable POSIX permissions for those users.

10 Select a user or group in the Access Control List field. Choose Allow or Deny from the Type field, and choose a level of access from the Permission field.

11 If you chose Custom in the Permission field, click the disclosure triangles to display specific attributes. Choose Allow or Deny from the Permission Type pop-up menu. Select specific permissions and click OK.

You can further grant or deny specific permissions that you cannot specify by only using POSIX permissions. For example, you can allow a user to list folder contents, but disallow that same user from reading file attributes.

12 Click Save.

Configuring AFP Share Points

If you are going to supply network home folders, you should use AFP, because it provides authentication-level access security. A user has to log in with a valid user name and password to access files.

You can also enable AFP using an SSH-secured tunnel for file sharing. This tunnel prevents intruders from intercepting your communication with an AFP share point. You cannot enable SSH-secured tunnels for AFP share points that host home folders.

For more information, see "Configuring AFP File Sharing Service" on page 237.

Configuring SMB/CIFS Share Points

You should not use SMB/CIFS unless you're hosting a share point specifically for Windows users. There are well-known risks associated with SMB/CIFS. For example, SMB/CIFS uses NTLMv1 and NTLMv2 encryption, both of which are very weak password hashing schemes.

For more information, see "Configuring Windows File Sharing Service" on page 241.

Configuring NFS Share Points

NFS file access is based not on user authentication, but on the user ID and the client IP address, so it is generally less secure than AFP. Use NFS only if you must provide home folders for a large number of users who use UNIX workstations.

Use Workgroup Manager to restrict access to an NFS share point, so that only required computers can access it.

To restrict access to an NFS share point:

1. Open Workgroup Manager.
2. Click Sharing.
3. Select a share point.
4. Click Protocols.
5. In the pop-up menu, choose NFS Export Settings.
6. If only a few computers need access to the share point, select "Export this item and its contents to" and choose Client in its pop-up menu. Click Add. Enter the IP address of a client computer.

   Only add client computers that require access to the share point.
7. If every computer in a subnet requires access to the share point, select "Export this item and its contents to" and choose Subnet. Enter the subnet address in the Subnet address field, and the subnet mask in the Subnet mask field.
8. Select "Map Root user to nobody." Select "Map All users to nobody."

   A user with "nobody" privileges has "Others" POSIX permissions.
9. Select "Read-only."
10. Click Save.

Configuring FTP Share Points

You should only enable FTP share points if you require anonymous access. Files are transferred from FTP share points unencrypted over the network. Transferring files over FTP does not guarantee confidentiality or file integrity.

If you need to use FTP for file transfers, use the `sftp` command. The `sftp` command provides a secure means of authentication and data transfer while FTP does not. For more information, see the `sftp` man page.

For more information about setting up FTP share points, see "Configuring FTP File Sharing Service" on page 238.

Controlling Network Views

Use network views to control what is seen by users of a particular computer when they click the Network icon in the sidebar of a Finder window, or when they choose Go > Network in the Finder.

Use the Network pane in Workgroup Manager to create and manage network views.

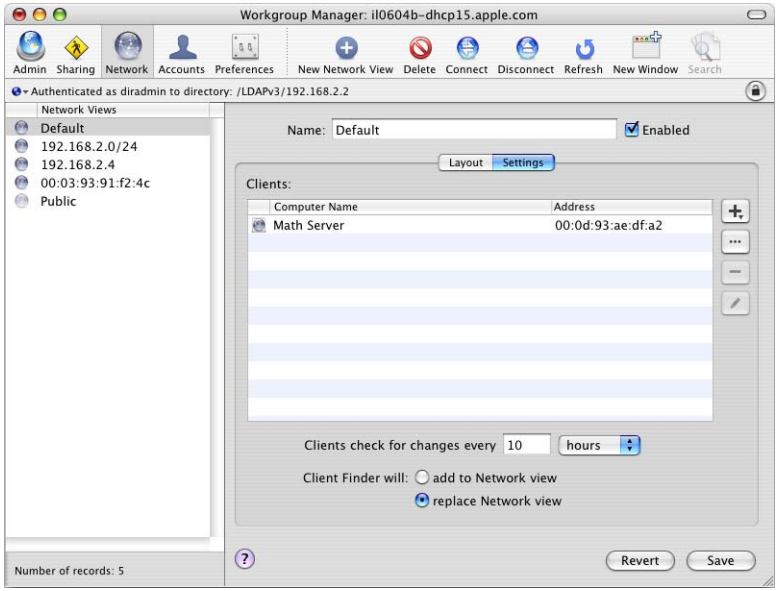

For more information about setting up network views, see the user management guide.

Use network views to limit the user's awareness of available servers to only those that the user must access. By lowering the awareness of available servers, unsophisticated intruders are not aware of possible networked computers. This also helps prevent accidental misuse of servers that your users might try to access.

Schools offer an example of the usefulness of network views. A school that has separate servers per department (one for English, one for math) should have a named view that only displays the math server. Named views are only visible on specific computers. All math labs in the school should then be configured so that this is the only network view that they use.

If a school has servers that it wants all students to be able to access, the school can use a default view, viewable by students with computers that have the school's directory domain on its search policy. The school can also create a public view that outsiders can view. This public view would have no visible servers, which would help prevent outsiders from becoming aware of available servers.

When you're creating network views, you should carefully define what resources are available in those views. For example, if you wanted to distribute antivirus software to students and, therefore, you include an IT software distribution server in a default view, you're exposing the server for a task that can be accomplished through using another method. You could distribute the antivirus software on an intranet site, where only students can download the software. This helps prevent the students from traversing your server.

You should not use dynamic lists like Bonjour to automatically populate your list. If you use dynamic list to populate your network view, you don't have as much control over what resources are included in the network view.

Securing Accounts

You can modify several account settings to improve security. Check with your organization to ensure that these settings do not conflict with network settings or organizational requirements.

In Workgroup Manager, you can use presets to save your settings as a template for future accounts. If you have settings that you apply to several accounts, you can use presets to expedite the creation of these accounts. Using presets also ensure that you use uniform account settings and help you avoid configuration errors. For more information, see the user management guide.

Configuring User Accounts

If you want to manage individual users or if you want those users to have unique identities on your network, create user accounts.

Before creating or modifying user accounts, you should have a firm understanding of what the account will be used for, and what authentication method you want to use.

To configure user accounts:

1. In Workgroup Manager, click Accounts.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Click Users, and select a user account. Click Basic.
4. If you want to grant server administration privileges to the user, select "administer the server."

   Server administration privileges allows the user to use Server Admin and make changes to a server's search policy using Directory Access.

5 If you want to grant directory domain administration privileges, select "administer this directory domain."

In the dialog that appears, click Users, Groups, or Computers to select the type of account you want to change privileges for. To remove editing privileges, deselect "Edit *account-type* preferences" and "Edit *account-type* accounts," where *account-type* is the type of account. To limit editing privileges to specific accounts, select "For *accounts* listed below," where *accounts* are users, groups, or computers. Drag accounts from the "Available users" list to the "For *accounts* listed below" list. Click OK.

By default, you are given full directory domain administration privileges.

6 Click Advanced. Deselect "Allow simultaneous login on managed computers."

By disallowing simultaneous login, you reduce the chances of version conflicts when loading and saving files. This helps remind users that they should log off of computers when they are not using them.

7 Choose the most secure password type available in the User Password Type pop-up menu.

If you don't use smart cards, you'll be able to choose either Open Directory or crypt password. Open Directory is a much more secure than crypt password. If your network uses Open Directory for authentication, you should authenticate with it. For more information about Open Directory and crypt passwords, see the Open Directory administration guide.

Smart cards are also a secure form of authentication. Smart cards use two-factor authentication, which helps ensure that your accounts are not compromised. For more information, see "Using Smart Cards" on page 74.

8 If you chose the Open Directory password type, click Options.

In the dialog that appears, select "Disable login on specific date" and enter the date that the user no longer needs the account. Select "Disable login after inactive for # days," and replace # with the number of days that indicates the user no longer needs the account. Select "Disable login after user makes # failed attempts," and replace # with 3.

Select "Allow the user to change the password." Select "Password must contain at least # characters," and replace # with 8. Select "Password must be reset every # days," and replace # with 90. If you want to require the user to create a new password during their next login, select "Password must be changed at next login."

Replace these suggested values with values that meet the requirements of your organization. Click OK.

9 Click Groups. Click the Add (+) button to open a drawer listing all available groups. Drag groups from the drawer into the Primary Group ID field or the Other Groups list.

A primary group is the group to which a user belongs to if the user does not belong to any other groups. If a user selects a different workgroup at login, the user still retains access permissions from the primary group.

The ID of the primary group is used by the file system when the user accesses a file he or she doesn't own. The file system checks the file's group permissions, and if the primary group ID of the user matches the ID of the group associated with the file, the user inherits group access permissions.

Adding a user to a group allows the user to access the group's group folder. Carefully choose which groups to add users to. For more information, see "Configuring Group Accounts" on page 138.

10 Click Home. Select a secure location for the user's home folder in the Home list. Enter an appropriate value in the Disk Quota field.

By using a disk quota, you prevent malicious users from performing a denial of service attack where they fill the home volume.

11 Click Mail. Select None.

If you must enable mail, select either POP only, or IMAP only, but not both. Using fewer protocols reduces the number of possible avenues of attack.

12 Click Info. Do not enter any information in the fields provided.

User information can be used by malicious attackers when they try to compromise the user's account.

13 Click Windows and then click Save.

Configuring Group Accounts

Create groups of individuals with similar access needs. For example, if you create a separate group for each office, you can specify that only members of a certain office can log in to certain computers. When you more specifically define groups, you have greater control over who can use what.

You can grant or deny POSIX or ACL permissions to groups. If you have nested groups, you can propagate ACL permissions to child groups.

Groups also have access to group folders, which provide an easy way for group members to share files with each other.

To configure group accounts:

1. In Workgroup Manager, click Accounts.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the name and password of a directory domain administrator.
3. Click Groups, and select a group account.
4. Click Members. Click the Add (+) button to open a drawer listing all users and groups. Drag users from the drawer into the Members list.

   Carefully choose which users you want to add to the group.
5. Click Group Folder. Select a secure location for the group folder in the Address list.
6. Click the Browse (...) button to open a drawer listing all users. Drag users from the drawer into the Owner Name field.

   The group folder owner is given read/write access to the group folder.
7. Click Save.

Configuring Computer Lists

A computer list comprises one or more computers that have the same preference settings and that are available to particular users and groups.

Every computer on your network should be a member of a computer list. If you don't assign a computer to a computer list, the computer is becomes a member of the Guest Computers computer list. By grouping computers under computer lists, you can restrict access to allow a few specific groups.

To configure computer lists:

1. In Workgroup Manager, click Accounts.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Click Computer Lists, and select a computer list you've created.

   The Guest Computers, Windows Computers, and All Computers computer lists are predefined groups. Assign computers to other computer lists so that you have more control over who can access them.
4. Click List. Add computers to the Computers list.

   Click the Add (+) button or Browse (...) button to add computers to the Computers list. Using the Add dialog requires that you know the computers ethernet ID. The Browse dialog allows you to select a computer from all computers found through Bonjour.

5 Click Access. Select "Restrict to groups below." Click the Add (+) button to open a drawer listing all groups. Drag groups from the drawer into the list.

You should only allow access to groups who must use the computers. You should update this list whenever groups no longer need to access the computers.

6 Deselect "Allow users with local-only accounts."

If you select "Allow users with local-only accounts," you should deselect "Local-only accounts pick workgroups from the above list," and select "Allow computer administrators to disable management."

If your network uses accounts stored in a network directory domain, you should not allow the use of local accounts. By allowing the use of local accounts you lose some control over who can use your computers. It is also harder to manage and maintain your network.

If you must allow local accounts, then those accounts should not be joining workgroups. Local computer administrators should be allowed to disable management, so that they can more easily perform local system diagnostics without having to circumvent computer management.

7 Click Cache. Enter a value other than 0 in the "Update the preferences cache every # *time_interval*" field.

Replace # with the amount of time, and choose a value of seconds, minutes, hours, days, or weeks for the *time_interval*.

Setting the amount of time to 0 turns off caching. Without caching, managed preferences do not take effect when the computer is disconnected from the network.

8 Click Save.

Managing Preferences

You can define preferences for user accounts, group accounts, and computer lists that are defined in a shared directory domain. A group with defined preferences is called a workgroup.

By managing preferences for users, workgroups, and computers, you can customize the user's experience and restrict users to accessing only the applications and network resources you choose. Properly set managed preferences help deter unsophisticated users from performing malicious activities. They can also help prevent users from accidentally misusing their computer.

Understanding Managed Preference Interaction

Your managed preferences interact differently depending on the preference you're managing and which account types you're applying the preference to.

The following illustration shows how managed preferences interact when the same preferences are set at multiple levels:

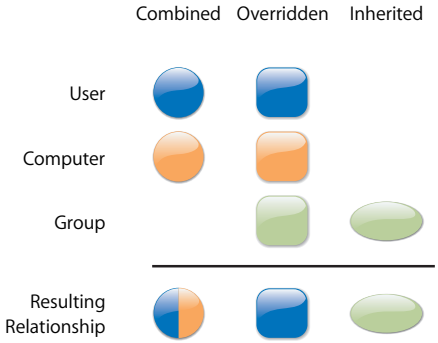

- Printing, Login, Applications, and some Dock preferences (involving items that appear in the Dock) are combined.

  For example, if you define printing preferences for users and computers, a user's printer list includes printers set up for both the user and the computer being used.

  Managed System Preferences are combined, in that different settings defined in Workgroup Manager act collectively at login.

- Other preference settings defined at more than one level can be overridden at login.

  The illustration below shows how overriding managed preferences interact when the same preferences are set at multiple levels:

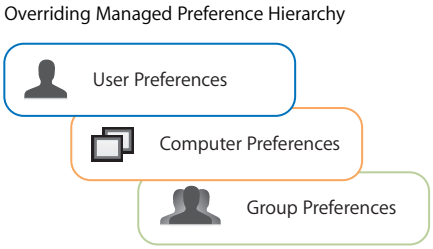

  When overriding preferences conflict, user preferences override both computer and group preferences, while computer preferences override group preferences.

For example, let's say you have different managed Dock preferences for users, workgroups, and computer lists. The Dock preferences for the user would take precedence, overriding and nullifying any Dock preferences set for workgroups or computers. If you do not manage any Dock preferences for the user, the computer list Dock preferences override and nullify any group Dock preferences.

For example, overriding preferences is useful in a school where you want to prevent all students from using recording devices attached to a school computer, except for students who serve as lab assistants. You can set up Media Access preferences for workgroups or computer lists to limit all students' access, but override these restrictions for lab assistants using Media Access settings at their user account level.

- Inherited preferences are preferences set at only one level.

  In some cases, you might find it easier and more useful to set certain preferences at only one level. In such a case, no overriding or combining occurs, and the user inherits the preferences without competition.

Choosing How to Manage Preferences

Most of the time you'll use workgroup-level and computer-level preferences.

- Workgroup preferences are most useful if you want to customize the work environment (such as application visibility) for specific groups of users, or if you want to use group folders.

  For example, a student might belong to a group named "Class of 2011" for administrative purposes and to a workgroup named "Students" to limit application choices and provide a group shared folder for turning in homework. Another workgroup might be "Teacher Prep," used to provide faculty members access to folders and applications for their use only.

- Computer-level preferences are useful when you want to manage preferences for users, regardless of their group associations. At the computer level, you might want to limit access to System Preferences, manage Energy Saver settings, list particular users in the login window, and prevent saving files and applications to recordable discs.

  Computer preferences also offer a way to manage preferences of users who don't have a network account, but who can log in to a Mac OS X computer using a local account. You'd set up a computer list that supports local-only accounts.

Setting the Permanence of Management

When you define preferences, you can choose to manage them Always or Once. They are set to Never by default.

- Always causes the preferences to remain in effect until you change them on the server. When properly designed, a Mac OS X application that conforms to standard preference conventions does not allow a user to modify preferences set to Always. You can use Always, for example, to make sure users can't add or remove Dock items.

  Some applications might allow the user to change the Always managed preference, but the next time the user logs back in, the preference reverts to the managed setting.

- Once is available for some preferences. You can create default preferences, which users can then modify and keep their modifications. These preferences are then effectively unmanaged. For example, you could set up a group of computers to display the Dock in a certain way the first time users log in. A user can change preferences you've set to Once, and the selected changes always apply to that user.

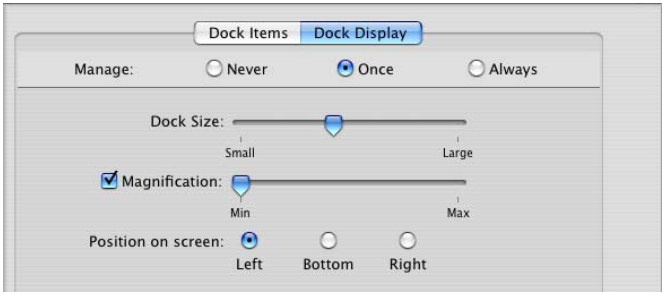

  In the Overview Preference panes, you can set the following preferences to Once: Dock, Finder (Preferences and Views), Internet, Login (Login Items), Mobility (Login & Logout Sync and Background Sync panes of Rules), and Universal Access. For all other preferences, you must choose either Always or Never.

- Never lets a user control his or her own preferences. However, some preference settings, such as Accounts and Date & Time, require a local administrator's user name and password before changes can be made. Never also means that the preferences are not managed at this account level, but can be managed at a different account level. For example, even if you set the Dock preference to Never for a particular user, the Dock preference could still be managed at the group or computer level.

Note: When using the preference editor (the Details view within the Preferences pane), you can set preferences to Often. Often settings are similar to Once settings, but are reapplied at every login. This management setting is useful for training environments. Users can customize their preferences to suit their needs during a session without any risk of affecting a future user's work experience. Additionally, some applications only respond to preference management if set to Often.

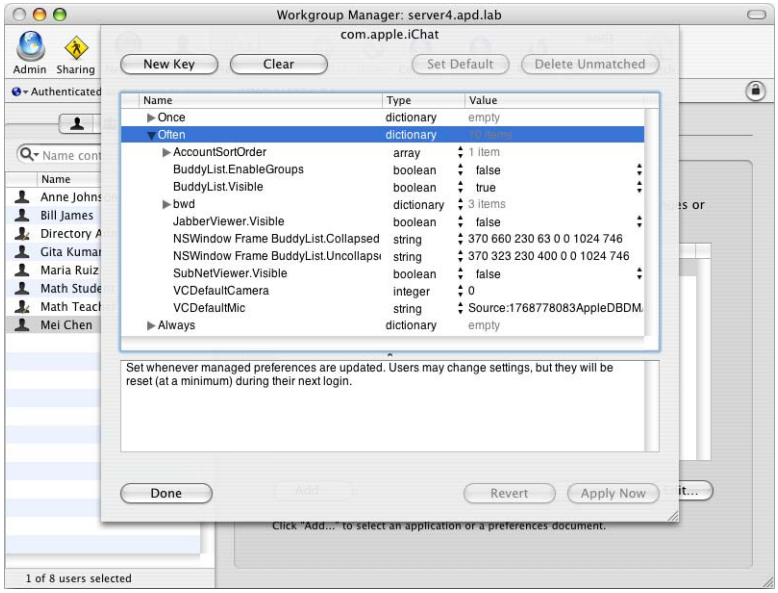

Managing Applications Preferences

Use settings in the Applications pane to provide users with access to applications. You can create lists of "approved" applications that users are allowed to open, and you can allow users to open items on local volumes. You can also prevent applications from opening restricted applications.

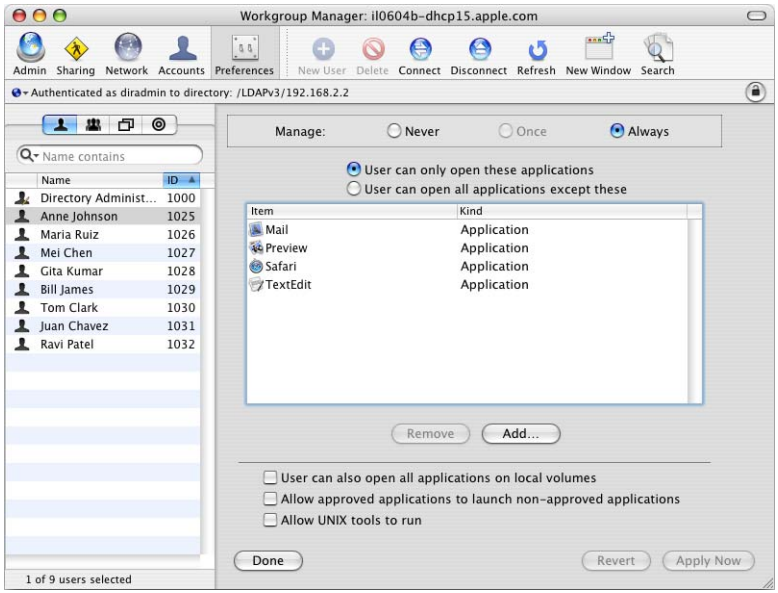

Note: Applications are identified by their bundle ID. Since a clever user might change an application's bundle ID and therefore defeat their access restrictions, the application restrictions should not be considered a barrier that no user can overcome.

To manage Applications preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Applications.
5. Select Always.
6. Select "User can only open these applications." To open a dialog where you can choose applications to add to the list, click Add. To remove applications from the list, select the application and click Remove.

   Modify this list to only include applications that the user requires, and your organization approves of.

Include any approved helper applications that approved applications might open. For example, if you give users access to an email application, you might also want to add a web browser, a PDF viewer, and a picture viewer to avoid problems opening and viewing email contents or attached files.

7 Deselect "User can also open all applications on local volumes."

If you enable this, users can access applications on the computer's local hard disk in addition to approved applications on CDs, DVDs, or other external disks.

8 Deselect "Allow approved applications to launch non-approved applications."

Instead of enabling this, add helper applications to the list.

9 Deselect "Allow UNIX tools to run."

If you choose not to allow access to these types of tools, some applications might not function properly.

10 Click Apply Now.

Managing Classic Preferences

Unless required, client computers should not run Classic. If your clients run Classic, there are a few management settings you can change to help secure it.

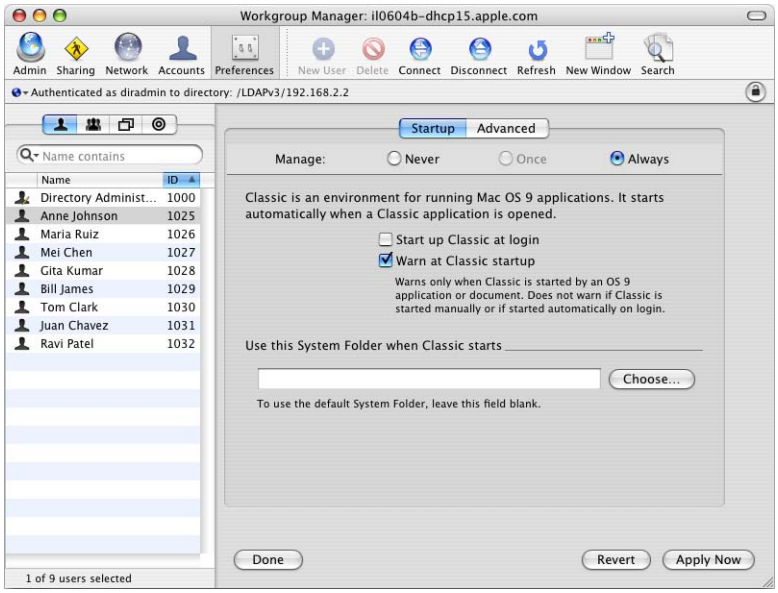

For more information about how to secure Classic locally, see "Securing Classic Preferences" on page 90.

To manage Classic preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Classic.
5. Click Startup. Select Always.
6. Deselect "Start up Classic at login."
7. Select "Warn at Classic startup."
8. Click Choose and select a Classic System Folder listed on a CD or DVD.
9. Click Advanced. Select Always.
10. Select "Allow special startup modes."
11. Click Apply Now.

Managing Dock Preferences

You can customize the user's Dock to display only certain applications. This helps you guide the user toward using specifically recommended applications.

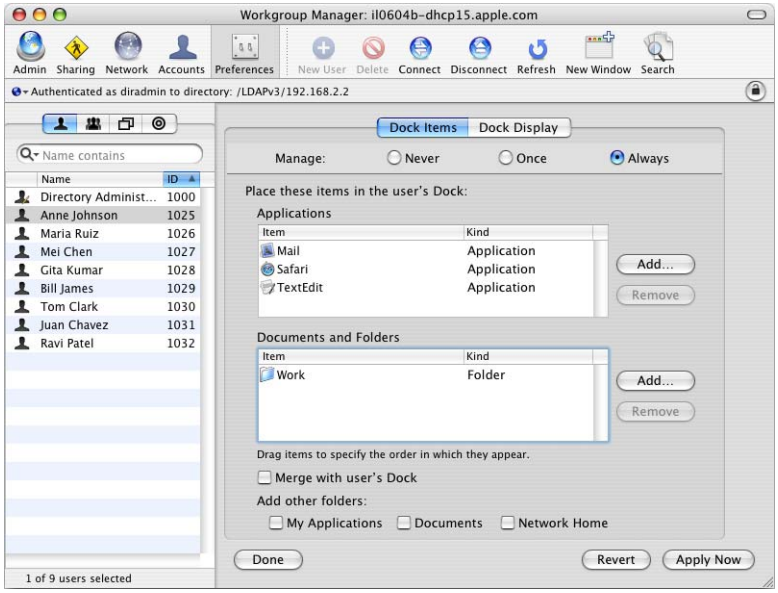

You can also add documents and folders to the Dock. Adding specific, required network folders to the Dock helps prevent the user from navigating through your network hierarchy. This also helps prevent them from accidentally misusing the server.

To manage Dock preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Dock.
5. Click Dock Items. Select Always.
6. In the Applications list, click Add to open a dialog where you can choose applications to add to the list. To remove applications from the list, select the application and click Remove.

   Modify this list to only include applications that the user requires and your organization approves of.
7. In the Documents and Folders list, click Add to open a dialog where you can choose documents and folders to add to the list. To remove documents and folders from the list, select the document or folder and click Remove.

   Modify this list to only include documents and folders that the user requires.
8. Deselect "Merge with user's Dock."

   By deselecting this, the user is not able to modify the Dock.
9. Deselect "My Applications," "Documents," and "Network Home."

   Limiting the number of items in your user's Dock helps guide users to required applications, documents, and folders.
10. Click Dock Display. Select Always.
11. Select "Automatically hide and show the Dock."

    This can help prevent others from seeing what applications your users have available when they casually pass by.
12. Click Apply Now.

Managing Energy Saver Preferences

Energy Saver preference settings help you save energy and battery power by managing wake, sleep, and restart timing for servers and client computers. You can only manage Energy Saver preferences for computer lists.

When client computers go to sleep, they become unmanaged. You should not enable sleep mode for any client computers.

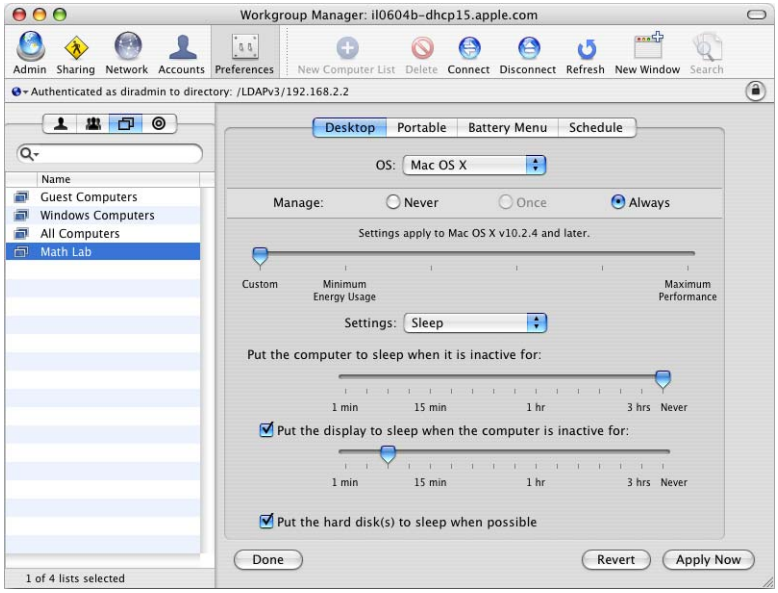

To manage Energy Saver preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select a computer list you created.

   You cannot set Energy Saver preferences for predefined computer lists.
4. Click Overview. Click Energy Saver.
5. Click Desktop. Choose Mac OS X from the OS pop-up menu. Select Always. Under "Put the computer to sleep when it is inactive for," move the slider to Never.
6. Choose Mac OS X Server from the OS pop-up menu. Select Always. Under "Put the computer to sleep when it is inactive for," move the slider to Never.
7. Click Portable. Choose Adapter from the Power Source pop-up menu. Select Always. Under "Put the computer to sleep when it is inactive for," move the slider to Never.

8 Choose Battery from the Power Source pop-up menu. Select Always. Under "Put the computer to sleep when it is inactive for," move the slider to Never.

9 Click Schedule. Select Always. Deselect "Start up the computer."

10 Click Apply Now.

Managing Finder Preferences

You can control various aspects of Finder menus and windows. By controlling Finder menus and windows, you can improve or control workflow.

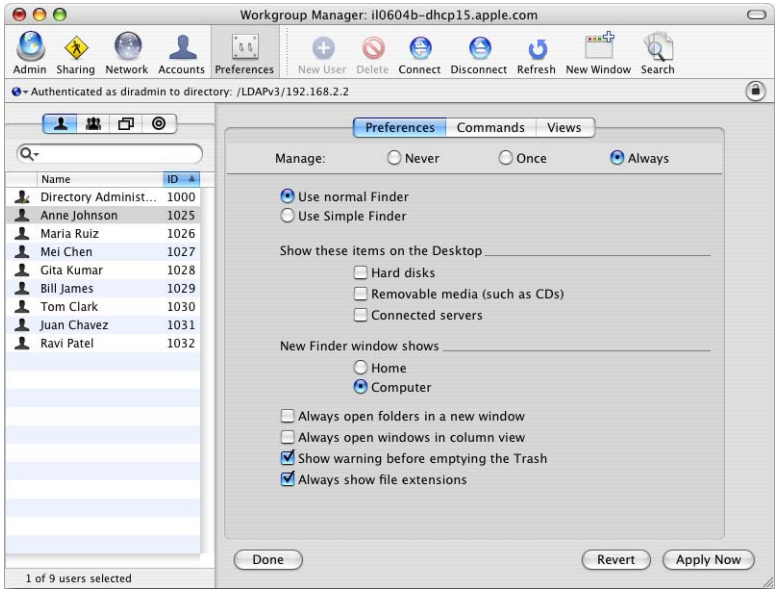

You can prevent users from burning media or from ejecting disks, and from connecting to remote servers. When used in conjunction with Dock preferences, you can guide the user experience.

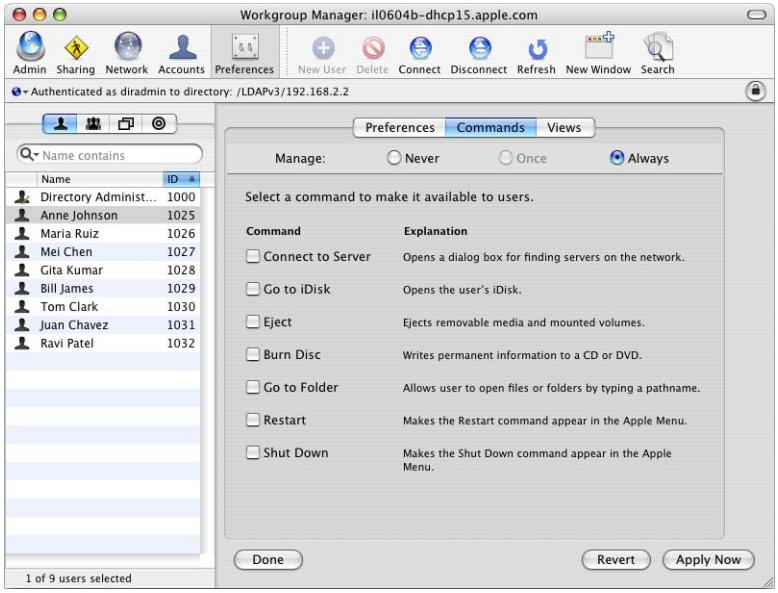

To manage Finder preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Finder.
5. Click Preferences (in the Finder Preferences pane). Select Always.
6. Select "Use normal Finder."

   Simple Finder is best used for computers in kiosk situations.

   Simple Finder removes the ability to use a Finder window to access applications or modify files. This limits users to accessing only what is in the Dock. If you enable Simple Finder, users cannot mount network volumes. With Simple Finder enabled, users cannot create folders or delete files.
7. Deselect "Hard disks," "Removable media (such as CDs)," and "Connected servers."

   By deselecting these, you help restrict novice users from casually navigating through local and network file systems.

8 Select "Always show file extensions."

Important: Operating systems use file extensions as one method of identifying types of files and their associated applications. Using only file extensions to check the safety of incoming files leaves your system vulnerable to attacks by Trojans. A Trojan is a malicious application which uses common file extensions or icons to masquerade as a document or media file (such as a PDF, MP3, or JPEG).

For further explanation and guidance on handling email attachments and content downloaded from the internet, see the KBase Article: #108009: Safety tips for handling email attachments and content downloaded from the Internet at: docs.info.apple.com/article.html?artnum=108009.

9 Click Commands. Select Always.

10 Deselect Connect to Server, Go to iDisk, and Go to Folder.

Instead of allowing the user to choose which servers or folders to load, you should add approved servers.

11 Deselect Eject and Burn Disc.

Disallowing external media gives you more control.

12 Deselect Restart and Shut Down.

By disallowing restarting and shutting down client computers, you help ensure that your computers are available to other users.

13 Click Apply Now.

Managing Internet Preferences

Internet preferences let you set email and web browser options. Some Internet browser or email applications might not support these settings.

By managing email preferences for users, you can reduce the chance of your users accidentally misconfiguring their email preferences.

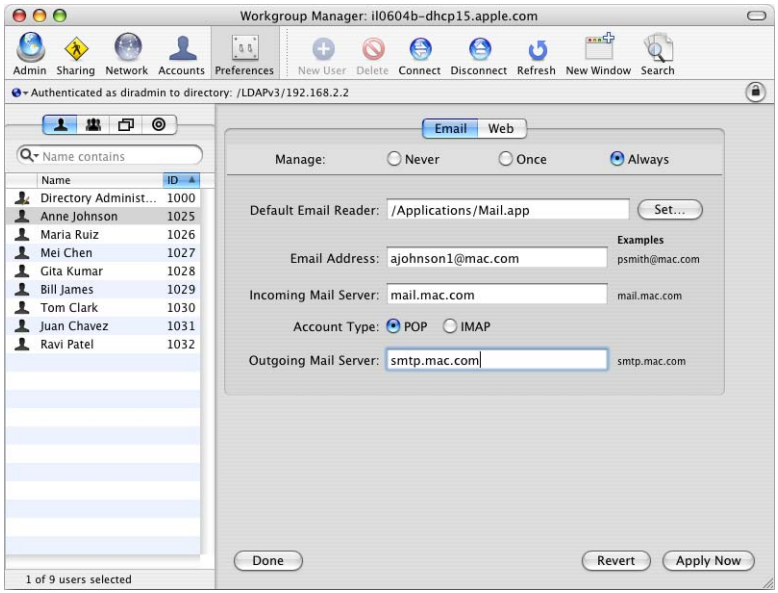

You can configure web preferences to direct your users to your organization's approved intranet and search page. This helps deter casual external access by your users.

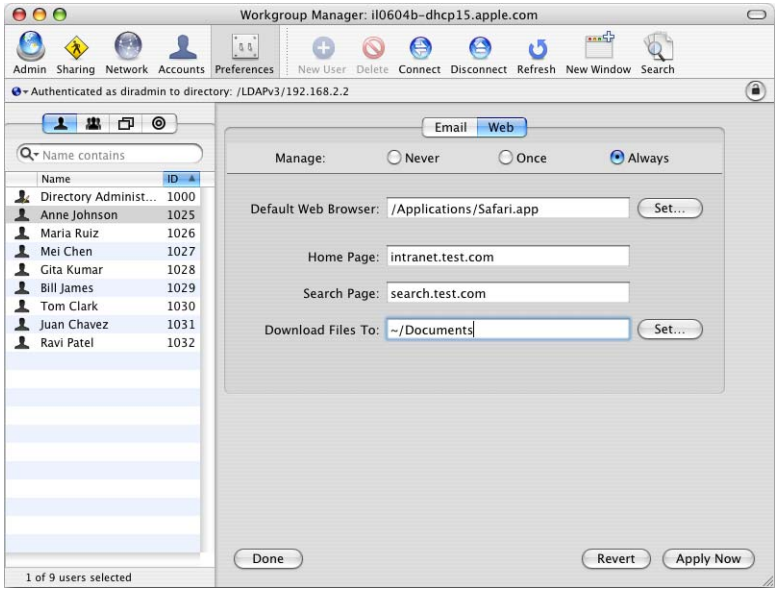

To manage Internet preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Internet.
5. Click Email. Select Always.
6. In the Default Email Reader field, click Set to open a dialog where you can choose an email application.
7. In Email Address, Incoming Mail Server, and Outgoing Mail Server, enter default user information. Select POP or IMAP.
8. Click Web. Click Always.
9. In the Default Web Browser field, click Set to open a dialog where you can choose an web browser application.
10. In Home Page and Search Page, enter intranet pages approved by your organization.
11. In the Download Files to field, click Set to open a dialog where you can choose a location in your user's home folder to download files.

    Choose a location within your user's home folder, not at the computer's root level.
12. Click Apply Now.

Managing Login Preferences

Use Login preferences to set options for user login, provide password hints, and control the user's ability to restart and shut down the computer from the login window. You can also mount a group volume or make applications open automatically when a user logs in.

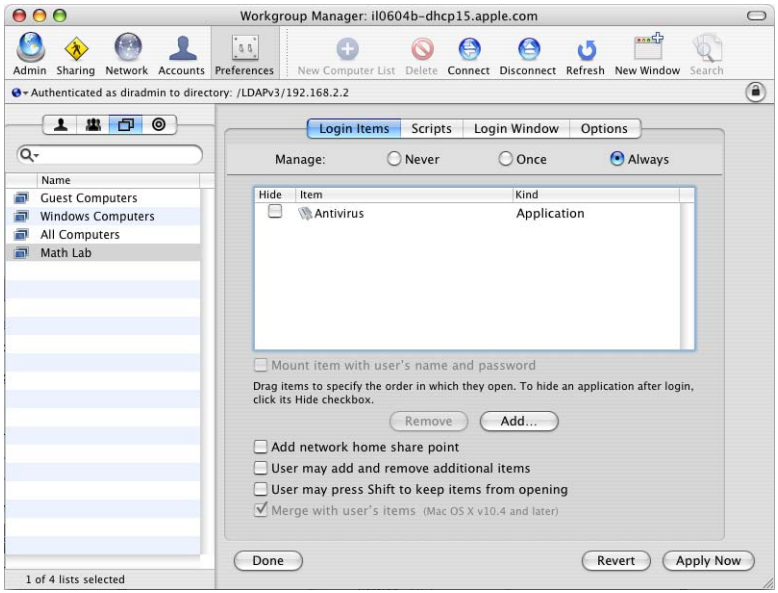

You can only apply script and login window settings to computer lists.

By managing script settings, you can help protect your users from malicious login or logout scripts that could be used to compromise their accounts integrity.

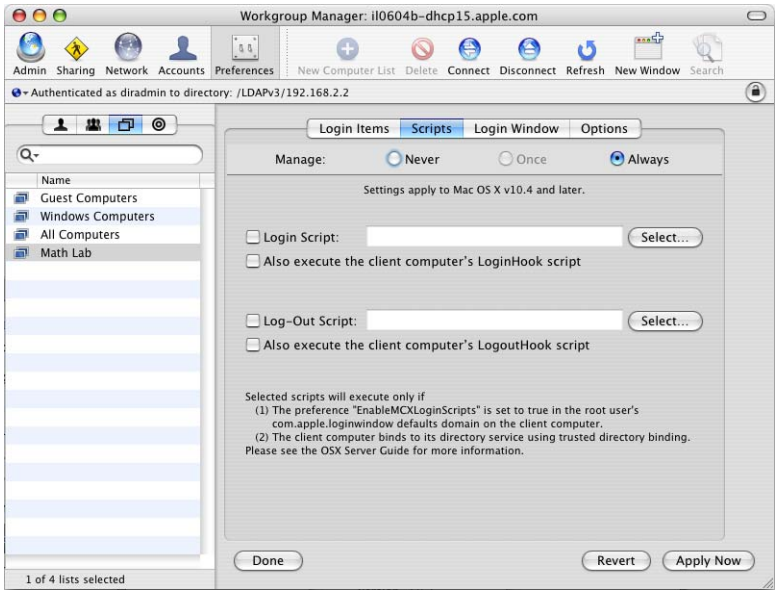

You can manage login window settings to make it more difficult for intruders to attempt to log in as legitimate users.

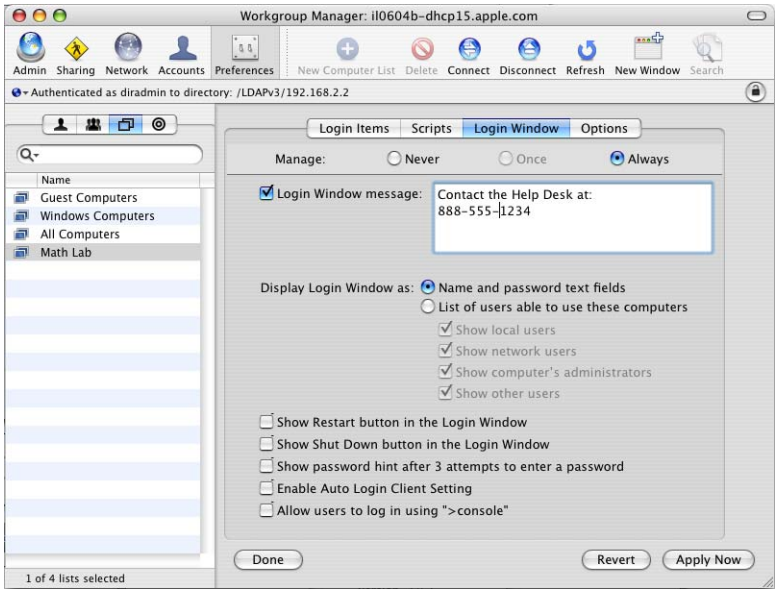

You can configure options so that you can track malicious user actions.

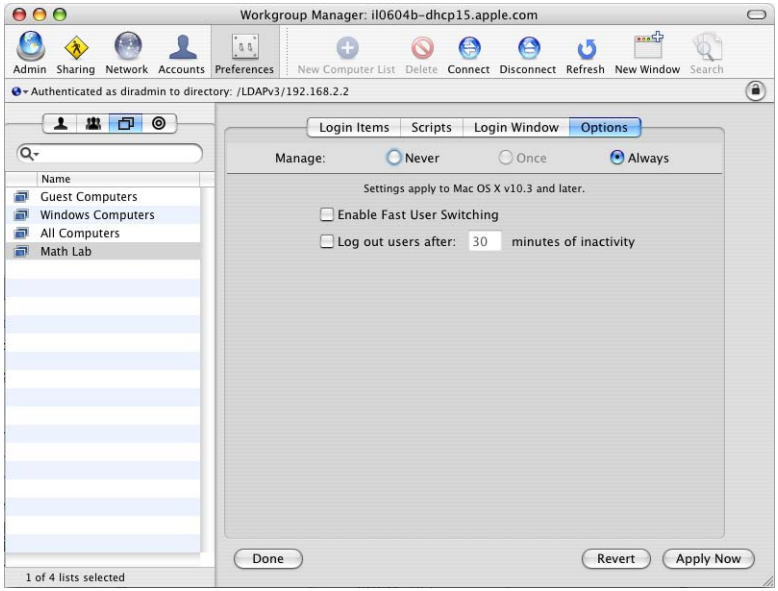

To manage Login preferences:

1 In Workgroup Manager, click Preferences.

2 Make sure the correct directory is selected and that you are authenticated for it.

To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.

3 Select an account.

To perform the steps involving applying scripts and login window settings, you must select a user-defined computer list.

4 Click Overview. Click Login.

5 Click Login Items. Select Always.

Different login items settings are available depending on whether you're managing Once or Always. Like all managed preferences, you should use the Always setting to ensure that your settings stay in effect past the user's first login.

6 To load applications or to mount a group volume at startup, click Add to open a dialog where you can add an application or volume.

Add antivirus and file integrity checking applications required by your organization.

7 Deselect "Add network home share point."

Instead of automatically mounting share points, the user should mount share points as required.

8 Deselect "User may add and remove additional items." Deselect "User may press Shift to keep items from opening."

Deselecting these options helps prevent the user from automatically loading possibly malicious applications. It also helps ensure that the user cannot bypass loading applications required by your organization.

9 Click Scripts. Select Always.

10 Unless your organization requires the use of specific login or logout scripts, deselect Login Script and Log-Out Script. Deselect "Also execute the client computer's LoginHook script," and "Also execute the client computer's LogoutHook script."

To run login and logout scripts, the client's computer has to achieve a certain level of "trust" with the server. This level of trust is based on how secure the client's connection is with the server. By requiring a certain level of trust, this ensures that the client computer does not run scripts from malicious servers.

For more information about how to enable the use of login and logout scripts, see the user management guide.

11 Click Login Window. Select Always.

12 Select "Login Window message" and enter help desk contact information in the adjacent field.

Do not enter any information about the computer's typical usage or who its users are.

13 In "Display Login Window as," select "Name and password text fields."

Requiring that users know their account names adds an additional layer of security and helps prevent unsophisticated intruders from compromising accounts with weak passwords.

14 Deselect "Show Restart button in the Login Window" and "Show Shut Down button in the Login Window."

Preventing users from easily restarting or shutting down the computer helps ensure that the computer is available to all users.

15 Deselect "Show password hint after 3 attempts to enter a password."

Password hints can help malicious users compromise accounts. If you enable this setting, set the password hint per user account to information for your organization's help desk.

16 Deselect "Auto Login Client Setting."

Enabling this setting allows users to enable automatic login through System Preferences. Automatic login bypasses all login window-based security mechanisms.

17 Deselect "Allow users to log in using '>console.'"

Enabling this setting allows the user to bypass the login window and use the Darwin console (command-line interface.)

18 Click Options. Select Always.

19 Deselect Enable Fast User Switching.

Fast User Switching allows multiple users to log in simultaneously. This makes it more difficult to track user actions and allows users to run malicious applications in the background while another user is actively using the computer.

20 Deselect "Log out users after # minutes of inactivity."

If you select "Log out users after # minutes of inactivity," make sure you enable password-protected screensavers in case a dialog prevents logging out.

21 Click Apply Now.

Managing Media Access Preferences

Media Access preferences let you control settings for, and access to, CDs, DVDs, the local hard disk, and external disks (for example, floppy disks and FireWire drives).

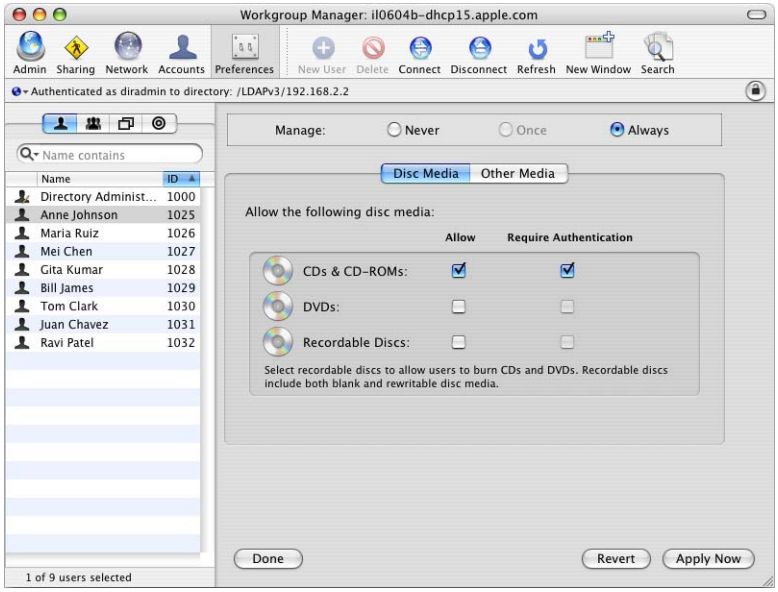

Disable all unnecessary media. If users can access external media, it provides more opportunities for performing malicious activities. For example, they can transfer malicious files from the media to the hard disk. Another example is if an intruder gains temporary access to the computer, he or she can quickly transfer confidential files to the media.

Carefully weigh the advantages and disadvantages of disabling certain forms of media. For example, disabling external disks prevents you from using USB flash memory drives for storing keychains. For more information, see "Storing Credentials" on page 75.

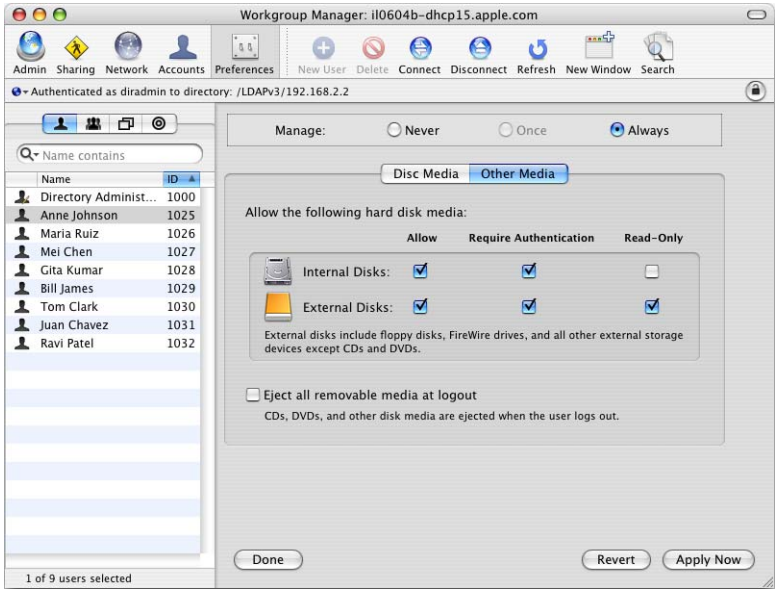

To manage Media Access preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.

   To perform the steps involving applying scripts and login window settings, you must select a user-defined computer list.
4. Click Overview. Click Media Access.
5. Select Always. Click Disc Media.
6. Unless you must use certain disc media, deselect Allow for CDs & CD-ROMs, DVDs, and Recordable Discs.

   If you need to enable certain disc media, select both Allow and Require Authentication for that disc media.
7. Click Other Media.

8 Unless you must use certain media, deselect Allow for Internal Disks and External Disks.

If you must enable certain media, select Allow and Require Authentication for that disc media. Select Read-Only if you do not need to save files to that media.

9 Select "Eject all removable media at logout."

This helps prevent users from forgetting they have media inserted in the computer.

10 Click Apply Now.

Managing Mobility Preferences

You can use Mobility preferences to enable and configure mobile accounts. Mobile accounts include both a network home folder and a local home folder. By having these two types of home folders, clients can take advantage of features available for both local and network accounts. You can synchronize specific folders of these two home folders, creating a portable home directory.

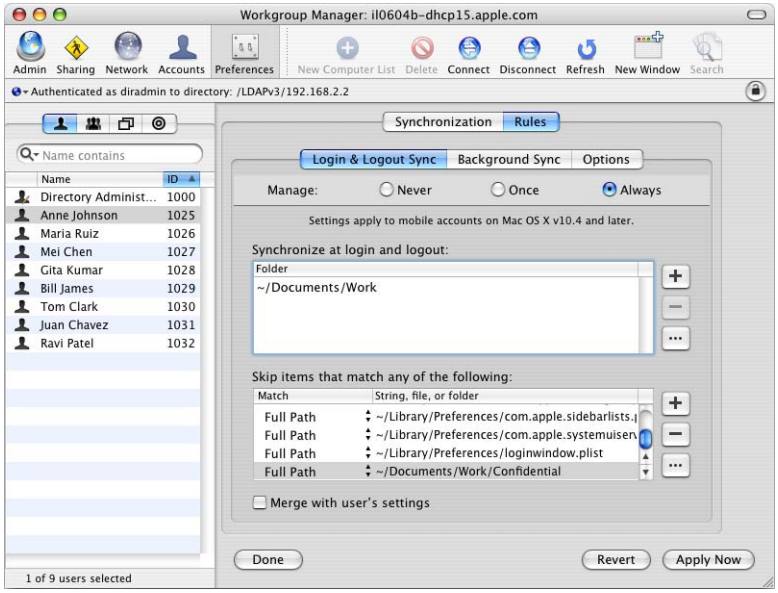

You shouldn't use mobile accounts. When you access a mobile account from a client computer and create a portable home directory, you create a local home folder on that client computer. If you access the mobile account from many computers, creating portable home directories on each of those computers, your home folder's files are stored on several computers. This provides additional avenues of attack.

If you use mobile accounts do not create portable home directories on any computers that are physically insecure, or that you infrequently access. Enable FileVault on every computer where you created portable home directories. For more information about enabling FileVault, see "Securing Security Preferences" on page 104.

To manage Mobility preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.

   To perform the steps involving applying scripts and login window settings, you must select a user-defined computer list.
4. Click Overview. Click Mobility.
5. Click Synchronization. Select Always.

   If you do not manage Synchronization settings, the user can choose to create his or her own mobile account in Accounts preferences.
6. To prevent the user from enabling a mobile account, deselect "Synchronize account for offline use." Click Apply Now, and do not continue following these instructions. If you select "Synchronize account for offline use," continue to follow these instructions.

   The rest of these instructions only apply if you decide to select "Synchronize account for offline use."
7. Select "Require confirmation before creating a mobile account."

   If deselected, a portable home directory is created every time the user accesses a different computer.
8. Click Rules. Click Login & Logout Sync. Select Always.
9. In the "Synchronize at login and logout" list, click the Add (+) button and enter the paths of folders located in the user's home folder. Alternatively, click the Browse (...) button to open a dialog where you can choose folders to add to the list.

   Add folders that do not contain confidential files.
10. In the "Skip items that match any of the following" list, click the Add (+) button and enter the paths of folders located in the user's home folder. Alternatively, click the Browse (...) button to open a dialog where you can choose folders to add to the list.

    Add folders that contain confidential files.
11. Deselect "Merge with user's settings."

    By deselecting this setting, the folders you choose to synchronize replace those chosen by the user.
12. Click Background Sync. Select Always.
13. In the "Synchronize at login and logout" list, click the Add (+) button and enter the paths of folders located in the user's home folder. Alternatively, click the Browse (...) button to open a dialog where you can choose folders to add to the list.

Add folders that do not contain confidential files.

14 In the "Skip items that match any of the following" list, click the Add (+) button and enter the paths of folders located in the user's home folder. Alternatively, click the Browse (...) button to open a dialog where you can choose folders to add to the list.

Add folders that contain confidential files.

15 Deselect "Merge with user's settings."

By deselecting this setting, the folders you choose to synchronize replace those chosen by the user.

16 Click Apply Now.

Managing Network Preferences

Network preferences let you select and configure proxy servers that can be used by users and groups. You can also specify hosts and domains to bypass proxy settings.

Using proxy servers controlled by your organization can help improve security. You can also decrease the performance hit from using proxies if you selectively bypass trusted hosts and domains (like choosing local resources or trusted sites).

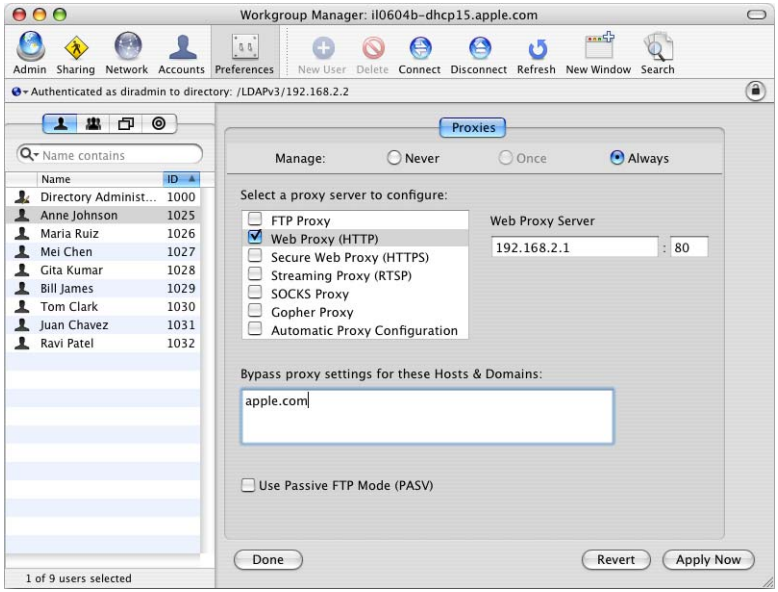

To manage Network preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.

   To perform the steps involving applying scripts and login window settings, you must select a user-defined computer list.
4. Click Overview. Click Network. Select Always.
5. Select a type of proxy server. Enter the network address and port of a proxy server controlled by your organization.
6. If you select Automatic Proxy Configuration, enter the URL of your automatic proxy configuration (.pac) file.
7. In the "Bypass proxy settings for these Hosts & Domains" field, enter the addresses of the hosts and domains that you want users to connect to directly. To enter multiple address, separate the subnet masks with new lines, spaces, semicolons, or commas. There are several ways to enter addresses:
   - A subdomain or fully qualified domain name (FQDN) of a target server, such as server1.apple.com or store.apple.com.
   - The specific IP address of a server, such as 192.168.2.1.
   - A domain name, such as apple.com. This bypasses apple.com, but not any subdomains, such as store.apple.com.
   - An entire website, including all subdomains, such as *.apple.com.
   - A subnet in Classless Inter-Domain Routing (CIDR) notation. For example, if you wanted to add a subnet of IP addresses from 192.168.2.0 to 192.168.2.255, you would name that view 192.168.2.0/24. For a detailed description of subnet masks and CIDR notation, see the network services administration guide.
8. Deselect Use Passive FTP Mode (PASV).
9. Click Apply Now.

Managing Printing Preferences

Printer preferences let you control which printers the user can access. Ideally, reduce the printer list to only those printers that the user needs to access.

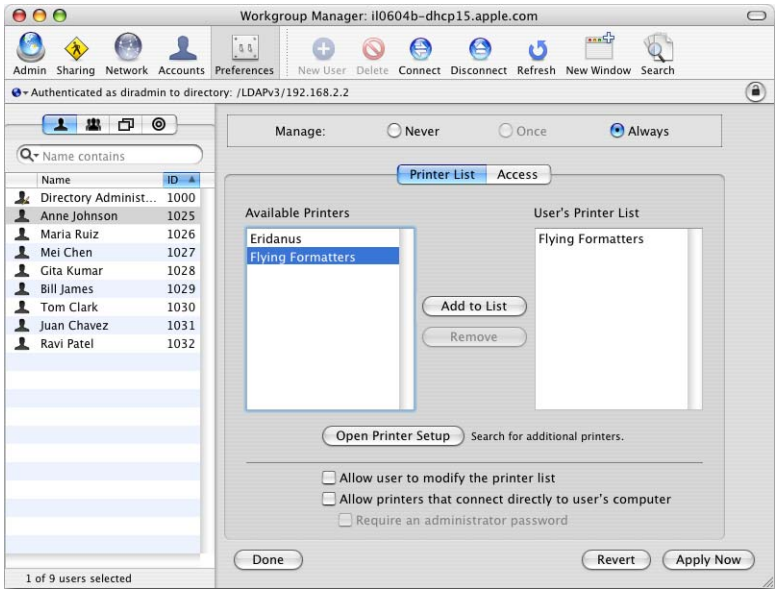

You should require that the user authenticate as an administrator before printing.

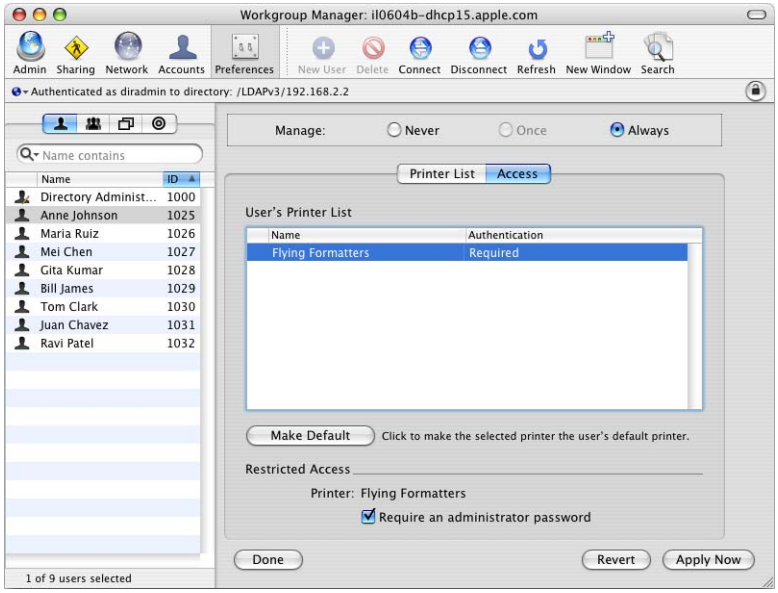

To manage Printing preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Printing. Select Always.
5. Select a printer in the Available Printers list and click Add to List.

   Add all printers that you want the user to access to the user's printer list.
6. If you want to add additional printers to the user's printer list, click Open Printer Setup.

   For more information, see Printer Setup Utility Help.
7. Deselect "Allow user to modify the printer list."
8. Deselect "Allow printers that connect directly to user's computer."

   If you select this setting, select "Require an administrator password."
9. Click Access.
10. Select a printer, and select "Require an administrator password."

    Repeat for all printers in the User's Printer List.
11. Click Apply Now.

Managing Software Update Preferences

With Mac OS X Server, you can create your own Software Update server to control the updates that are applied to specific users or groups. This is advantageous because it reduces external network traffic while also providing you more control. By configuring the Software Update server, you can choose which updates to provide.

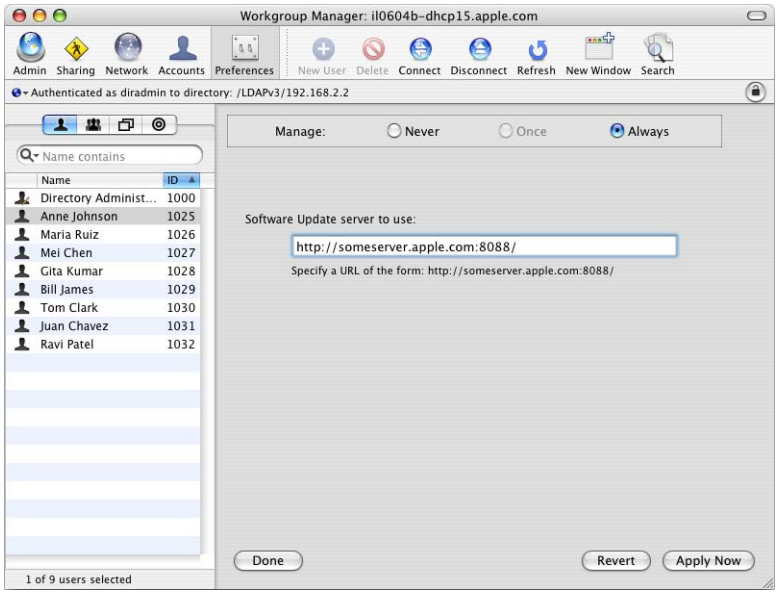

To manage Software Update preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click Software Update. Select Always.
5. Specify a URL of the form of http://someserver.apple.com:8088/.

   Enter the URL of your internal Software Update server.
6. Click Apply Now.

Managing System Preferences Preferences

You can specify which preferences are displayed in System Preferences preferences. If a user can display a particular preference, it does not necessarily mean that the user can modify that preference. Some preferences, such as Startup Disk preferences, require an administrator name and password before a user can modify its settings.

The preferences that appear in Workgroup Manager are those installed on the computer you're currently using. If your administrator computer is missing any preferences that you would like to disable on client computers, you should either install the applications related to those preferences or use Workgroup Manager on a computer that includes those preferences.

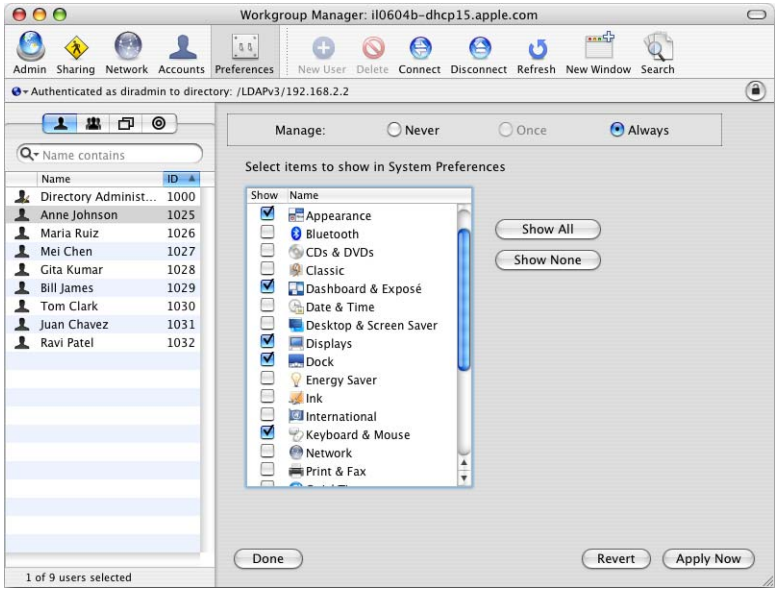

To manage System Preferences preferences:

1. In Workgroup Manager, click Preferences.
2. Make sure the correct directory is selected and that you are authenticated for it.

   To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.
3. Select an account.
4. Click Overview. Click System Preferences. Select Always.
5. Click Show None.
6. Select Appearance.
7. Select Dashboard & Expose.
8. Select Displays.

9 Select Dock.

10 Select Keyboard & Mouse.

11 Select Security.

12 Select Universal Access.

13 Click Apply Now.

Disabling Widgets

You can disable Dashboard for a network managed user, group, or computer list. The Dashboard widgets included with Mac OS X Server can be trusted. However, you should be careful when you install third-party Dashboard widgets. You can install Dashboard widgets without having to authenticate. Disabling Dashboard prevents unauthorized use of widgets.

To disable Dashboard widgets for a network managed user, group, or computer list:

1 Open Workgroup Manager and select a user, group, or computer list.

2 Select Preferences, then Applications.

3 Deselect the "Allow approved applications to launch non-approved applications" option to disable sublaunching.

4 This allows Dashboard to display, but unable to open any widgets. Any currently open widgets will close automatically.

Managing Universal Access Preferences

Universal Access settings can help improve the user experience for certain users. For example, if a user has a disability, has difficulty using a computer, or wants to work in a different way, you can choose settings that enable the user to work more effectively.

Most of the Universal Access settings do not negatively impact security. However, some settings allow other users to more easily see what you're doing.

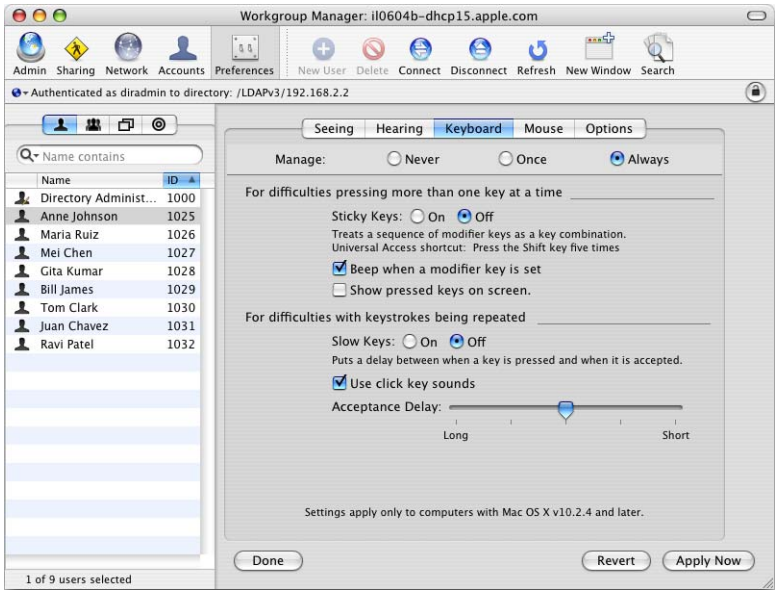

To manage Universal Access preferences:

1 In Workgroup Manager, click Preferences.

2 Make sure the correct directory is selected and that you are authenticated for it.

To switch directories, click the globe. If you are not authenticated, click the lock and enter the user name and password of a directory domain administrator.

3 Select an account.

4 Click Overview. Click Universal Access.

5 Click Seeing. Select Always.

6 Deselect Turn on Zoom.

Pressing and holding the Option, Command, and + keys will zoom in, while pressing and holding the Option, Command, and - keys will zoom out.

7 Click Keyboard. Select Always.

8 Select Sticky Keys Off. Deselect "Show pressed keys on screen."

If Sticky Keys are on and you select "Show pressed keys on screen," modifier keys, such as Control, Option, Command, and Shift are displayed on screen. All other keys are not displayed.

9 Click Apply Now.

Managing Certificates

8

Mac OS X Server supports many services which ensure encrypted data transfer. This encryption is facilitated through certificates.

Mac OS X Server uses a Public Key Infrastructure system to generate and maintain certificates of identities. Server Admin makes it easy to manage SSL certificates that can be used by web, mail, Open Directory, and other services that support them. You can create a self-signed certificate, and generate a Certificate Signing Request (CSR) to obtain an SSL certificate from an issuing authority and install the certificate.

For more information about how to use SSL certificates with individual services, see Chapter 9, "Setting General Protocols and Access to Services," on page 183.

Understanding Public Key Infrastructure

Public Key Infrastructure (PKI) systems allow the two parties in a data transaction to be authenticated to each other, and to use encryption keys and other information in identity certificates to encrypt and decrypt messages traveling between them.

PKI enables multiple communicating parties to establish confidentiality, message integrity, and message source authentication without having to exchange any secret information in advance.

Secure Sockets Layer (SSL) technology relies on a PKI system for secure data transmission, and user authentication. It creates an initial secure communication channel to negotiate a faster, secret key transmission. Mac OS X Server uses SSL to provide data encrypted data transmission for mail, web, and directory services.

The following sections contain more background information about key aspects of PKI:

- Public and Private Keys
- Certificates
- Certificate Authorities (CA)
- Identities

Public and Private Keys

Within a PKI, two digital keys are created: the public key and the private key. These keys are mathematically linked such that data encrypted with one key can only be decrypted by the other, and vice versa. If a user named Tom publicly distributed his public key, then user Anne could use it to encrypt a message and send it to him. Only Tom is able to decrypt and read the message because only he has his private key.

In this scenario, Anne still has to verify the key that is supposedly from Tom is really from him. Suppose a malicious user posing as Tom sent Anne his own public key. The malicious user would then be able to decrypt Anne's message, which might have been intended for Tom only.

To verify that it's really Tom who is sending Anne his public key, a trusted third party can verify the authenticity of Tom's public key. In SSL parlance, this trusted third party is known as a certificate authority (CA). The CA signs Tom's public key with its private key, creating a certificate. Now, anyone can verify the certificate's authenticity using the CA's public key.

The private key isn't meant to be distributed to anyone, and often is itself encrypted by a passphrase. The public key, on the other hand, is distributed to other communicating parties. Basic key capabilities can be summed up as:

Key Type	Capabilities
Public keys	• Can encrypt messages that can only by decrypted by the holder of the corresponding Private key. • Can verify the signature on a message originating as coming from a Private key.
Private keys	• Can digitally sign a message or certificate, claiming authenticity. • Can decrypt messages which were encrypted with the Public key. • Can encrypt messages which can only be decrypted by the Private key itself.

Web, Mail, and Directory Services use the public key with SSL to negotiate a shared key for the duration of the connection. For example, a mail server sends its public key to a connecting client and initiates negotiation for a secure connection. The connecting client uses the public key to encrypt a response to the negotiation. The mail server, since it has the private key, can decrypt the response. The negotiation continues until both the mail server and the client have a shared secret to encrypt traffic between the two computers.

Certificates

Public keys are often contained in certificates. A user can digitally sign messages using his or her private key, and another user can verify the signature using the public key contained in the signer's certificate, which was issued by a CA within the PKI.

A public key certificate (sometimes called an "identity certificate") is a file in a specified format (Mac OS X Server uses the x.509 format) which contains:

- The public key half of a public-private key pair.
- The key user's identity information, such as a person's user name and contact information.
- A validity period (how long the certificate can be trusted to be accurate).
- The URL of someone with the power to revoke the certificate (its "revocation center").
- The digital signature of either a CA, or the key user himself.

Certificate Authorities

A certificate authority (CA) is an entity that signs and issues digital identity certificates claiming trust of the identified party. In this sense, it's a trusted third party between two transactions.

In x.509 systems, CAs are hierarchical in nature, with CAs being certified by CAs, until you reach a "root authority." The hierarchy of certificates is always a top-down, with a root authority's certificate at the top. A root authority is a CA that's trusted by enough or all of the interested parties, so that it doesn't need to be authenticated by yet another trusted third party.

A CA can be a company that, for a fee, signs and issues a public key certificate that states the CA attests that the public key contained in the certificate belongs to its owner, as recorded in the certificate. In a sense, CA is a "digital notary public." One applies to the CA for a certificate by providing identity and contact information, as well as the public key. A CA must check an applicant's identity, so that users can trust certificates issued by that CA to belong to the identified applicant.

Identities

Identities, in the context of the Mac OS X Server Certificate Manager, are the combination of a signed certificate for both keys of a PKI key pair. The identities are used by the system keychain, and are available for use by various services that support SSL.

Self-Signed Certificates

Self-signed certificates are certificates that are digitally signed by the private key of the key pair included in the certificate. This is done in place of a CA signing the certificate. By self-signing a certificate, you're attesting that you are who you say you are. No trusted third party is involved.

Readying Certificates

Before you can use SSL in Mac OS X Server's services, the certificates must be created or imported. You can create your own self-signed certificate, generate a Certificate Signing Request (CSR) to send to a CA, or import a certificate previously created with OpenSSL.

Using Certificate Manager

Certificate Manager is integrated into Server Admin to help you create, use, and maintain identities for SSL-enabled services. Certificate Manager provides integrated management of SSL certificates in Mac OS X Server for all services that allow the use of SSL certificates.

Certificate Manager allows creation of self-signed certificates and certificate signing requests (CSRs) to obtain a certificate signed by a CA. The certificates, either self-signed, or signed by a CA, are accessible by the services that support SSL.

Identities that were previously created and stored in SSL files can also be imported into Certificate Manager, where they are accessible to all the services that support SSL.

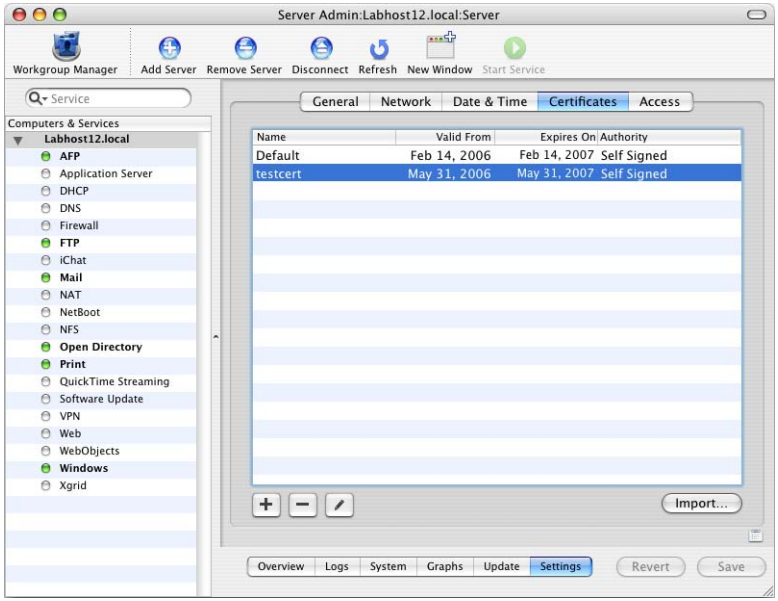

Certificate Manager displays the following for each certificate:

- The domain name for which the certificate was issued.
- Its dates of validity.
- Its signing authority, such as the CA entity. If the certificate is self-signed, it reads "Self-Signed."

Certificate Manager in Server Admin does not allow you to sign and issue certificates as a CA, nor does it allow you to sign and issue certificates as a root authority. If you need any of these functions, you can use Certificate Assistant, located in /System/Library/CoreServices/. Certificate Assistant allows these functions and others. Self-signed and CA-issued certificates created using Certificate Assistant can be used in Certificate Manager by importing the certificate.

Requesting a Certificate from a CA

Certificate Manager helps you create a certificate signing request (CSR) to send to your designated CA.

To request a signed certificate:

1. Open Server Admin.
2. In the Computers & Services list, select the server you are requesting a certificate for.
3. Click Settings.
4. Click Certificates.
5. Click the Add (+) button.
6. Fill out identity information.

   The common name is the fully qualified domain name of the server which uses SSL enabled services.
7. Enter starting and ending validity dates.
8. Select a private key size (1024 bits is the default).
9. Enter a passphrase for the private key and Click Save.

   This passphrase should be more secure than a normal password. You should use at least 20 characters, including mixed case, numbers, and punctuation. Don't repeat characters and don't use words contained in the dictionary.
10. Click "Request Signed Certificate. . ."
11. Follow the onscreen directions for requesting a signed certificate from your chosen CA.
12. Click Send Request.
13. Click Done.
14. When the CA replies to the email, the signed certificate is included in the email text.
15. Click Add Signed Certificate.
16. From your CA certificate email, copy the characters from "==Begin CSR==" to "==End CSR==" into the text box.
17. Click OK.
18. Click Save.

Creating a Self-Signed Certificate

Whenever you create an identity in Certificate Manager, you're creating a self-signed certificate. Certificate Manager creates a public-private key pair in the system keychain with the key size specified (512–2048 bits). It then creates the corresponding self-signed certificate in the system keychain.

A Certificate Signing Request (CSR) is also generated at the same time that the self-signed certificate is created. This isn't stored in the keychain, but is written to disk at /etc/certificates/cert.common.name.tld.csr, where "common.name.tld" is the Common Name of the certificate that was issued.

To create a self-signed certificate:

1. Open Server Admin.
2. In the Computers & Services list, select the server you are creating a certificate for.
3. Click Settings.
4. Click Certificates.
5. Click the Add (+) button.
6. Fill out identity information.

   The common name is the fully qualified domain name of the server which uses SSL enabled services.
7. Enter starting and ending validity dates.
8. Select a private key size (1024 bits is the default).
9. Enter a passphrase for the private key and Click Save.

   This passphrase should be more secure than a normal password. You should use at least 20 characters, including mixed case, numbers, and punctuation. Don't repeat characters and don't use words contained in the dictionary.
10. Click Save.

Importing a Certificate

You can import a previously generated SSL certificate and private key into Certificate Manager. The items are stored in the list of identities and are available to SSL-enabled services.

To import an existing SSL certificate:

1. Open Server Admin.
2. In the Computers & Services list, select the server you are importing a certificate to.
3. Click Settings.
4. Click Certificates
5. Click Import.

6 In the Certificate File field, enter the existing certificate's file name and path.

   Alternately, click Browse and locate your certificate file.

7 In the Private Key File field, enter the existing private key file's name and path.

   Alternately, click Browse and locate your private key file.

8 In the Certificate Authority File field, enter the existing certificate authority file's name and path.

   Alternately, click Browse and locate your certificate authority file.

9 Enter the private key passphrase.

10 Click Import.

Modifying Certificates

Once a certificate is created and signed, you shouldn't have to do much more with the certificates. They are only editable in Server Admin and cannot be changed once a CA signs the certificate. Self-signed certificates can be changed. Certificates should be deleted if the information they possess (such as contact information) is no longer correct or if you believe the key pair has been compromised in some way.

Editing a Certificate

Once the certificate signature of a CA is added, it can't be edited. A self-signed certificate can be edited. All the fields of the certificate (including domain name and private key passphrase, private key size, and so on) can be modified. If the identity was exported to disk from the system keychain, it will have to be reexported.

To edit a certificate:

1 Open Server Admin.

2 In the Computers & Services list, select the server with the certificate you are editing.

3 Click Settings.

4 Click Certificates.

5 Select the Certificate Identity to edit.

   It must be a self-signed certificate.

6 Click the Edit (/) button.

7 Modify the certificate settings.

8 Click Save.

Deleting a Certificate

When a certificate has expired, or been compromised, you'll need to delete it.

To delete a certificate:

1. Open Server Admin.
2. In the Computers & Services list, select the server with the certificate you are deleting.
3. Click Settings.
4. Click Certificates.
5. Select the Certificate Identity to delete.
6. Click the Delete (–) button to delete the certificate.
7. Click Save.

Creating a Certificate Authority

If your server must communicate using SSL with external computers out of your control, you should purchase SSL certificates from a well-known CA. Once the certificates have been obtained, configuration of the services is the same whether they were purchased from a vendor or signed by your own CA.

If you are setting up an internal network and only need to encrypt local traffic, set up a CA to sign SSL certificates for the internal network. The next sections describe this process. While the security is only as good as the security of the CA, in many cases this is sufficient to enable encrypted communication between a web or mail server and their clients. The basic steps to set up an internal SSL-encrypted network are:

- Create a CA.
- Distribute the CA's certificate to client systems.
- Use the CA to sign the certificates the servers will use.

Using Certificate Assistant

Mac OS X Server includes the Certificate Assistant application, which allows you to sign and issue certificates as a CA, and allows you to sign and issue certificates as a root authority. Certificate Assistant is located in /System/Library/CoreServices/ and is available as a menu item from Keychain Access.

Creating a CA Using Certificate Assistant

Since the security of your certificates is dependent on the security of the CA, performing these steps on a secure computer is critical. The computer should be physically secure and not connected to any network.

To create the CA using Certificate Assistant:

1. Open Certificate Assistant and click Continue.
2. Select "Create a Certificate Authority (CA)."
3. Deselect "Certificate will be 'self-signed' (root)."

   Selecting this option creates a self-signed root certificate authority that are often used for testing purposes in place of certificates signed by proper CAs.
4. Fill out certificate information.

   The common name is the fully qualified domain name of the server which uses SSL enabled services.

   The validity period is the number of days the certificate is valid. Click Continue.
5. Choose an issuer for the certificate. An issuer signs the certificate you are going to create. Click Continue.
6. Select the key size (2048 bits is the default) and algorithm (RSA is the default) used to create your key pair for the CA. Click Continue.
7. Select the key size (2048 bits is the default) and algorithm (RSA is the default) that specify the public and private key pair information for users of this CA when they request a certificate. Click Continue.
8. Set the Key Usage Extension (KUE) for this CA.

   The key usage extensions of a CA's certificate identify its security capabilities and determines how the certificate can be used. For example a certificate can be created to sign emails, but not encrypt them.

   Deselect "This extension is critical" if it is safe for the software using the certificate to ignore the extension if unrecognized. Otherwise, if software does not recognize a critical extension, the certificate is rejected. Click Continue.
9. Set the Key Usage Extension for users of this CA.

   Certificate key usage extensions are also set for the users of this CA if required. Click Continue.
10. Set the miscellaneous extensions for this CA by selecting "Include Basic Constraints extension (Extension is always critical)" and "Use this certificate as a certificate authority." Click Continue.

    The basic constraint extension indicates whether the certificate is a CA and the maximum allowable depth of the certificate chain. Select "Include Subject Alternate Name extension" for this CA, if required.

    This allows the CA to use additional names for the certificate subject and provides for flexible controls.

11 Set the miscellaneous extensions for the users of this CA by selecting "Include Basic Constraints extension (Extension is always critical)" and "Use this certificate as a certificate authority."

Miscellaneous extensions are also set for the users of this CA if required.

12 Select "Include Subject Alternate Name extension," if required for the users of the CA. Click Continue.

13 Specify a location for the certificate by choosing a keychain where the certificate will be stored. Click Continue.

14 Create a CA configuration file by entering the name of the CA configuration file. This file can be used by others to easily request a certificate from you.

15 Select "Make this CA the default," if necessary. Click Continue.

Your CA is then created and is ready to issue certificates.

Creating a CA from the Command Line

Since the security of your certificates is dependent on the security of the CA, performing these steps on a secure computer is critical. The computer should be physically secure and not connected to any network.

To create the CA using the openssl command:

1 Enter the following in Terminal to create a certificate directory.

```
$ cd /usr/share
$ sudo mkdir certs
$ cd certs
```

2 Generate a key pair with the `openssl` command.

```
$ sudo openssl genrsa -des3 -out ca.key 2048
```

This command generates a Triple-DES encrypted RSA public-private key pair names ca.key. The `2048` is the length of the key in bits. OpenSSL asks for a passphrase for the key upon creating it. Use a strong passphrase and keep it secure. A compromise of this passphrase undermines the security of your entire certificate system.

Signing a Newly Created CA

After the key pair, the public key is signed to create an SSL certificate that can be distributed to other systems. Later, when we sign other servers' certificates with our CA's private key, any client can then use the CA's SSL certificate (containing its public key) to verify those signatures. When a CA signs a server's certificate with its private key, it means that it is vouching for the authenticity of those certificates. Anyone who can trust the CA can then trust any certificate the CA signs.

To sign the newly created CA's public key to produce a certificate for distribution:

```
$ sudo openssl req -new -x509 -days 365 -key ca.key -out ca.crt
```

When prompted, enter a strong passphrase for the key, as well as these fields:

```
Country Name:
Organizational Unit:
State or Province Name:
Common Name:
Locality Name (city):
Email Address:
Organization Name:
```

These fields should be filled out as accurately as possible, but those that don't apply can be left blank. At least one field must be filled in.

This creates a self-signed certificate named ca.crt, using the keys in ca.key, which is valid for a year (365 days). This limit can be set to a longer period of time, although this is less secure. The issue is similar to changing passwords regularly. A balance must be found between convenience and security.

Storing the CA Private Key

The CA private key should be generated on a computer that is not connected to your internal network. For added security, you can store the keychain containing the private key on USB storage so that you can keep the CA private key unavailable when connected to the network.

Creating Folders and Files for SSL

When signing certificates, SSL looks for keys and related information in directories specified in its configuration file openssl.cnf, which is found in /System/Library/OpenSSL/.

To create the directories and files where SSL expects to find them by default:

```
$ cd /usr/share/certs
$ sudo -s
$ mkdir -p demoCA/private
$ cp ca.key demoCA/private/cakey.pem
$ cp ca.crt demoCA/cacert.pem
$ mkdir demoCA/newcerts
$ touch demoCA/index.txt
$ echo "01" > demoCA/serial
$ exit
```

Now the CA is ready to sign certificates for servers, enabling encrypted communication between servers and clients.

Deploying Server Certificates to Clients

If you're using self-signed certificates, a warning pops up in most user applications saying that the certificate authority is not recognized. Other software, such as the LDAP client, simply refuses to use SSL if the server's CA is unknown. Mac OS X Server ships only with certificates from well-known commercial CAs. To prevent this warning, your CA certificate must be exported to every client computer that will be connecting to the secure server.

To distribute the self-signed CA certificate:

1. Copy the self-signed CA certificate (the file named ca.crt) onto each client computer.

   This is preferably distributed using non-rewritable media, such as a CD-R. Using non-rewritable media prevents the certificate from being corrupted.

2. Double-click the ca.crt icon where it was copied onto the client computer, to open the Keychain Access tool.

3. Add the certificate to the X509Anchors keychain using Keychain Access.

   Alternatively, use the `certtool` command in Terminal:

   ```
   $ sudo certtool i ca.crt k=/System/Library/Keychains/X509Anchors
   ```

   Now, any client application that checks against the system's X509Anchors keychain (such as Safari and Mail) recognizes any certificate signed by your CA.

Setting General Protocols and Access to Services

9

Use Server Admin to configure access to services and set general protocols.

Server Admin helps you configure and manage your servers. Using Server Admin, you can set general protocols, name or rename computers, set the date and time, manage certificates, and set user access to specific services.

Setting General Protocols

Mac OS X Server includes basic network management protocols, including network time protocol (NTP) and simple network management protocol (SNMP). Unless these are required, they should be disabled.

Disabling NTP and SNMP

The NTP software allows computers on a network to synchronize their Date & Time settings. Client computers specify their NTP server in the Date & Time panel of System Preferences. If NTP service is required, enable it on a single, trusted server on the local network. This service should be disabled on all other servers. For more information about the open source implementation, visit www.ntp.org.

The SNMP software allows other computers to monitor and collect data on the state of a computer running Mac OS X Server. This helps administrators identify computers that warrant attention, but use of this service is not recommended.

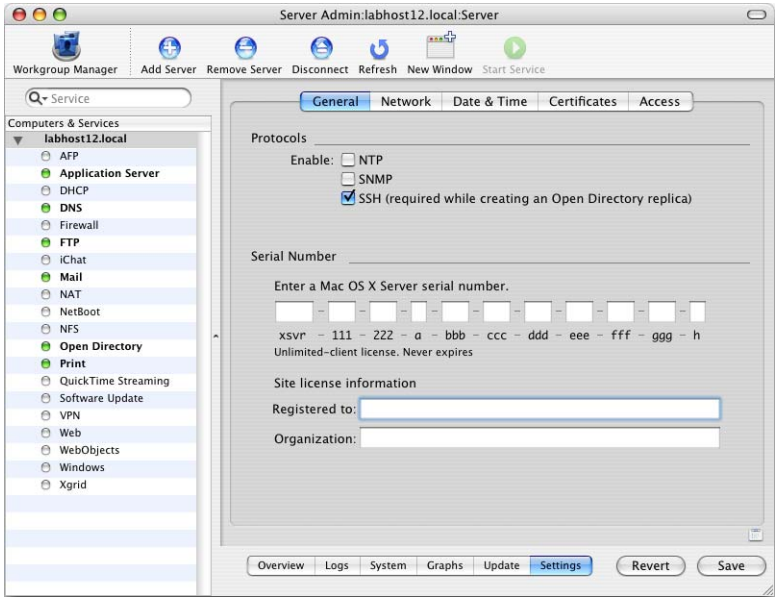

To disable NTP and SNMP:

1. Open Server Admin.
2. Click the name of the server you are configuring.
3. Click Settings.
4. Click General.
5. Deselect "Enable NTP" and "Enable SNMP."
6. Click Save.

Enabling SSH

Mac OS X Server also includes secure shell (SSH). The SSH software allows you to log in to other computers on a network, execute commands remotely, and move files from one computer to another. It provides strong authentication and secure communication, and is therefore recommended. For more information, see www.openssh.org.

To enable SSH:

1. Open Server Admin.
2. Click the name of the server you are configuring.
3. Click Settings.
4. Click General.
5. Select "Enable SSH (required while creating an Open Directory replica)."
6. Click Save.

Setting the Server's Host Name

You can change your computer name and local host name in Server Admin. When other users use Bonjour to discover your available services, the server is displayed as *hostname*.local.

To increase your privacy, you should change the host name of your computer so that your computer cannot be easily identified. The name should not indicate the purpose of the computer and the word "server" should not be used as the name or part of the name.

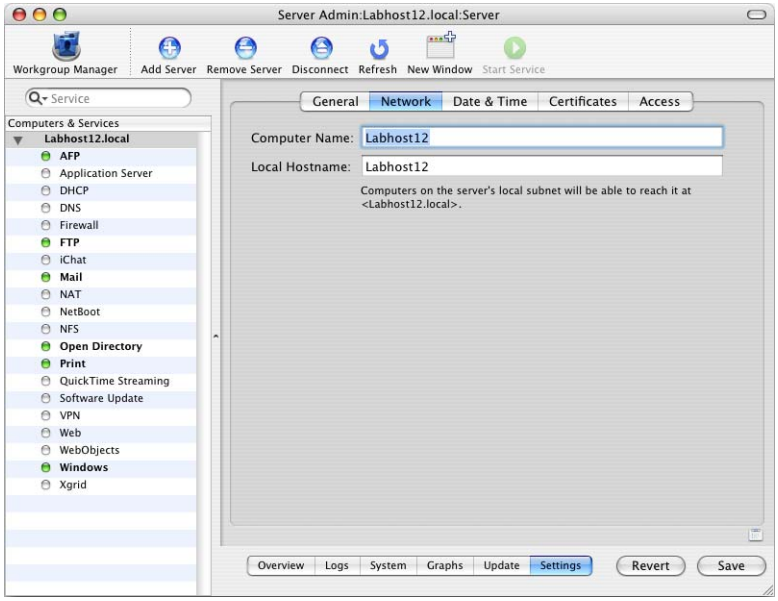

Setting the Date and Time

Correct date and time settings are required for authentication protocols, like Kerberos. Incorrect date and time settings can cause security issues. You can use Server Admin to configure your computer to automatically set the date and time based on an NTP server. If you require automatic date and time, use a trusted, internal NTP server.

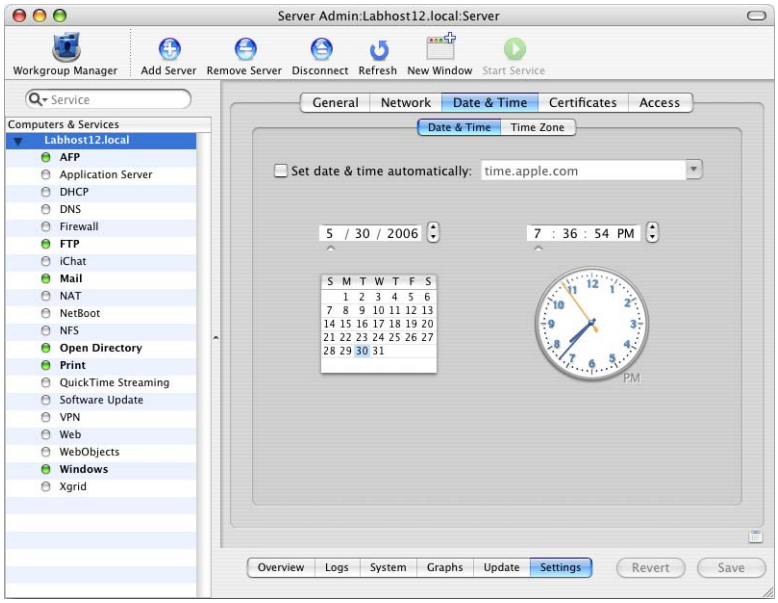

Setting Up Certificates

Certificate Manager is integrated into Server Admin to help you create, use, and maintain identities for SSL-enabled services. Certificate Manager provides integrated management of SSL certificates in Mac OS X Server for all services that allow the use of SSL certificates.

For more information about setting up certificates, see “Readying Certificates” on page 174.

Setting Service Access Privileges

You should use service access control lists (SACLs) so that services are only available to specific users. SACLs let you specify which users and groups have access to AFP, FTP, and Windows file services. SACLs allows you to add another layer of access control on top of standard and ACL permissions. Only users and groups listed in a SACL have access to its corresponding service. For example, if you want to prevent users from accessing AFP share points on a server, including home folders, remove the users from the AFP service's SACL.

Open Directory authenticates user accounts, and SACLs authorize use of services. If Open Directory authenticates you, the SACL for login window determines whether you can log in, the SACL for AFP service determines whether you can connect for file service, and so on. Some services also determine whether a user is authorized to access particular resources. This authorization may require retrieving additional user account information from the directory domain. For example, AFP service needs the user ID and group membership information to determine which folders and files the user is authorized to read and write.

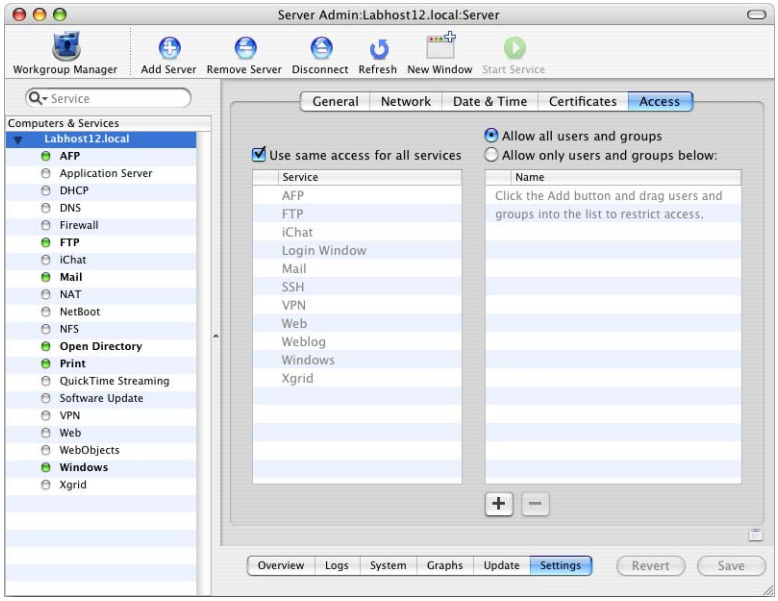

To set SACL permissions for a file service:

1. Open Server Admin.
2. Select the server in the Computers & Services list.
3. Click Settings.
4. Click Access.
5. Deselect "Use same access for all services" and select a service from the Service list. This will allow you to customize access to each service as required.

   To restrict access to all services, select "Use same access for all services."
6. If you want to restrict access to certain users and groups, select "Allow only users and groups below."
7. Click the Add (+) button to open the Users & Groups drawer.
8. Drag users and groups from the Users & Groups drawer to the list.
9. Click Save.

You can limit access to command-line tools that might run services, by limiting the use of the `sudo` command. For more information, see "Restricting sudo Usage" on page 71.

Securing Remote Access Services

10

Many organizations have individuals who need to connect to network resources remotely. This can create additional vulnerabilities unless your remote access services are securely configured.

Mac OS X Server allows remote access using remote login and VPN services. These services should be disabled unless they are required.

Remote Access services via remote login consists of two components each using the Secure Shell (SSH) service to establish an encrypted tunnel between the client and server. "Securing Remote Login" on page 191 discusses securing the server component, while "Configuring Secure Shell" on page 192 discusses securing the client component.

For additional information about configuring remote access services, see the network services administration guide.

Securing Remote Login

You can use SSH to remotely log in to Mac OS X Server. SSH creates a secure encrypted channel that protects communication with your computers. Older services that do not encrypt their communications, such as Telnet or RSH, should never be used—they allow network eavesdroppers to intercept passwords or other data.

Unless you must remotely log in to the computer or use another program that depends on SSH, disable the remote login service. However, Server Admin requires that you enable SSH. If you disable remote login, you cannot use Server Admin to remotely administer the server.

To disable remote login:

1. Open System Preferences.
2. Click Sharing.
3. Deselect Remote Login in the Service list.

Configuring Secure Shell

SSH lets you send secure, encrypted commands to a remote computer, as if you were sitting at the computer. Use the `ssh` tool in Terminal to open a command-line connection to a remote computer. While the connection is open, commands you enter are performed on the remote computer.

Note: You can use any application that supports SSH to connect to a computer running Mac OS X or Mac OS X Server.

SSH works by setting up encrypted tunnels using public and private keys. Here is a description of an SSH session:

1. The local and remote computers exchange their public keys. If the local computer has never encountered a given public key before, both SSH and a web browser prompt you whether to accept the unknown key.
2. The two computers use the public keys to negotiate a session key that is used to encrypt all subsequent session data.
3. The remote computer attempts to authenticate the local computer using RSA or DSA certificates. If this is not possible, the local computer is prompted for a standard user-name/password combination. See "Generating Key Pairs for Key-Based SSH Connections" on page 194 for information about setting up certificate authentication.
4. After successful authentication, the session begins. Either a remote shell, a secure file transfer, a remote command, or so on, is begun through the encrypted tunnel.

You should be aware of the following SSH tools:

- `sshd`—Daemon that acts as a server to all other commands
- `ssh`—Primary user tool: remote shell, remote command, and port-forwarding sessions
- `scp`—Secure copy, a tool for automated file transfers
- `sftp`—Secure FTP, a replacement for FTP

Modifying the SSH Configuration File

If you need to use SSH, you should alter the default settings. The SSH server configuration file is /private/etc/sshd_config.

To enable SSH, see "Enabling SSH" on page 185.

To modify the SSH configuration file:

1. Open the /private/etc/sshd_config file from Terminal.

```
$ sudo pico /private/etc/sshd_config
```

Authenticate, if requested.

This loads the sshd_config file in the `pico` text editor. For more information, see the `pico` man page.

2 Locate the "Authentication" section.

3 Disable root login using SSH. Replace the `PermitRootLogin` line with:

```
PermitRootLogin no
```

This forces the administrator to use `su` or `sudo` to obtain root privileges.

4 Verify permissions. Replace the `StrictModes` line with:

```
StrictModes yes
```

This has the SSH server verify that users' file and folder permissions are correct before allowing the connection.

5 Allow access for specific users (for example user1, user2, and user3) by adding the following line to the file:

```
AllowUsers user1 user2 user3
```

By default, SSH allows normal user accounts to log in.

6 Deny access for specific users (for example user4, user5, and user6) by adding the following line to the file:

```
DenyUsers user4 user5 user6
```

7 Deny access to users not using SSH version 2. Verify that the following line exists in your installation:

```
Protocol 2
```

By default, the configuration file specifies that only version 2 of the SSH protocol is supported. Using only version 2 is strongly recommended.

The following table includes additional security settings you can use when configuring the sshd_config file for your organization.

SSH options	Description
#PermitRootLogin	Allows or prevents logging in as root through SSH
#PasswordAuthentication	Enables or disables password authentication
#PermitEmptyPasswords	Permits or denies access to accounts without passwords
#PubKeyAuthentication	Enables or disables key-based authentication
#RSAAuthentication	Enables or disables RSA authentication (not needed for key-based authentication)
#RhostsRSAAuthentication	Enables or disables Rhost authentication (not needed for key-based authentication)
#ChallengeResponseAuthentication	Specifies whether challenge response authentication is allowed (not needed for key-based authentication)
#UsePAM	Enables or Disables PAM authentication and session set up. If you enable this, you should probably disable PasswordAuthentication (not needed for key-based authentication)
#StrictModes	Ensures that files and folders are adequately protected by the server's permissions' scheme

SSH options	Description
#LoginGraceTime	Changes the time allowed to authenticate
#KeyRegenerationInterval	Time interval that the server key is changed
#ServerKeyBits	Specifies the server key bit length
#Protocol	Specifies which SSH protocol to accept or reject
#Banner	Sends a warning message from a specified file to the user before authentication

For more information, see the `sshd_config` man page.

Generating Key Pairs for Key-Based SSH Connections

By default, SSH supports the use of password, key, and Kerberos authentication. The standard method of SSH authentication is supplying login credentials in the form of a user name and password. Identity key pair authentication enables you to log in to the server without having to supply a password.

This process works as follows:

1. A private and a public key are generated, each associated with a user name to establish that user's authenticity. When you attempt to log in as that user, the user name is sent to the remote computer.
2. The remote computer looks in the user's .ssh/ folder for the user's public key. This folder is created after using SSH the first time.
3. A challenge is then sent to the user based on his or her public key.
4. The user verifies his or her identity by using the private portion of the key pair to decode the challenge.
5. Once decoded, the user is logged in without the need for a password. This is especially useful when automating remote scripts.

Key-based authentication is more secure than password authentication because it requires that you have the private key file and know the password that lets you access that key file. Password authentication can be compromised without needing a private key file.

Note: If the server uses FileVault to encrypt the home folder of the user you want to use SSH to connect as, you have to be logged in on the server to be able to use SSH. Alternatively, you can store the keys for the user in a location that is not protected by FileVault. However, this is not very secure.

To generate the identity key pair:

1. Enter the following command on the local computer.

```
$ ssh-keygen -t dsa
```

2. When prompted, enter a filename to save the keys in the user's folder in.

3 Enter a password followed by password verification (empty for no password). For example:

```
Generating public/private dsa key pair.
Enter file in which to save the key (/Users/anne/.ssh/id_dsa): frog
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in frog.
Your public key has been saved in frog.pub.
The key fingerprint is:
4a:5c:6e:9f:3e:35:8b:e5:c9:5a:ac:00:e6:b8:d7:96 annejohnson1@mac.com
```

This creates two files. Your identification or private key is saved in one file (*frog* in our example) and your public key is saved in the other (*frog.pub* in our example). The key fingerprint, which is derived cryptographically from the public key value, is also displayed. This secures the public key, making it computationally infeasible for duplication.

Note: The location of the server SSH key is /etc/ssh_host_key.pub. Back up your key in case you need to reinstall your server software. If your server software is reinstalled, you can retain the server identity by putting the key back in its folder.

4 Copy the resultant public file, which contains the local computer's public key, to the .ssh/ folder in the user's home folder on the remote computer. The next time you log in to the remote computer from the local computer, you won't need to enter a password.

Note: If you are using an Open Directory user account and have already logged in using the account, you do not have to supply a password for SSH login. On Mac OS X Server computers, SSH uses Kerberos for single sign-on authentication with any user account that has an Open Directory password (Kerberos must be running on the Open Directory server). See the Open Directory administration guide for more information.

Updating SSH Key Fingerprints

The first time you connect to a remote computer using SSH, the local computer prompts for permission to add the remote computer's fingerprint (or encrypted public key) to a list of known remote computers. You might see a message like this:

```
The authenticity of host "server1.example.com" can't be established.
RSA key fingerprint is a8:0d:27:63:74:f1:ad:bd:6a:e4:0d:a3:47:a8:f7.
Are you sure you want to continue connecting (yes/no)?
```

The first time you connect, you have no way of knowing whether this is the correct host key. Most people respond "yes." The host key is then inserted into the ~/.ssh/known_hosts file so it can be compared against in later sessions. Be sure this is the correct key before accepting it. If at all possible, provide your users with the encryption key either through FTP, email, or a download from the web, so they can be sure of the identity of the server.

If you later see a warning message about a man-in-the-middle attack when you try to connect, it might be because the key on the remote computer no longer matches the key stored on the local computer. This can happen if you:

- Change your SSH configuration on either the local or remote computer.
- Perform a clean installation of the server software on the computer you are attempting to log in to using SSH.
- Start up from a Mac OS X Server CD on the computer you are attempting to log in to using SSH.
- Are attempting to SSH in to a computer that has the same IP address as a computer that you previously used SSH with on another network.

To connect again, delete the entries corresponding to the remote computer you are accessing (which can be stored by both name and IP address) in ~/.ssh/known_hosts.

Important: Removing an entry from the known_hosts file bypasses a security mechanism that would help you avoid imposters and man-in-the-middle attacks. Be sure you understand why the key on the remote computer has changed before you delete its entry from the known_hosts file.

Controlling Access to SSH

You can use Server Admin to control which users can open a command-line connection using the `ssh` tool in Terminal. Users with administrator privileges are always allowed to open a connection using SSH. The `ssh` tool uses the SSH service.

For information about restricting user access to services, see "Setting Service Access Privileges" on page 188.

Understanding SSH Man-in-the-Middle Attacks

An attacker might be able to get access to your network and compromise routing information, such that packets intended for a remote computer are instead routed to the attacker who impersonates the remote computer to the local computer and the local computer to the remote computer.

Here's a typical scenario: a user connects to the remote computer using SSH. By means of spoofing techniques, the attacker poses as the remote computer and receives the information from the local computer. The attacker then relays the information to the intended remote computer, receives a response, and then relays the remote computer's response to the local computer. Throughout the process, the attacker is privy to all the information that goes back and forth, and can modify it.

The following message may indicate a man-in-the-middle attack when connecting to the remote computer using SSH.

```
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@    WARNING: REMOTE HOST IDENTIFICATION HAS CHANGED!     @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
```

Protect against this type of attack by verifying that the host key sent back is the correct host key for the computer you are trying to reach. Be watchful for the warning message, and alert your users to its meaning.

Transferring Files Using SFTP

SFTP is a secure FTP protocol that uses SSH to transfer files. SFTP encrypts both commands and data, preventing passwords and sensitive information from being transmitted over the network. SFTP should always be used instead of FTP.

To transfer a file using SFTP:

1. Open Terminal.
2. Start the SFTP session.

   ```
   $ sftp username@hostname
   ```

   Where *`username`* is your user name and *`hostname`* is the IP address or host name of the server you are connecting to.
3. Enter your password when prompted.
4. You are now connected securely to the server and can use the SFTP commands to transfer files from the prompt.

   ```
   sftp>
   ```

   Use the `put` command to transfer a file from the local computer to the remote computer. Use the `get` command to transfer a file from the remote computer to the local computer.
5. Enter the following to transfer a picture file from the remote computer to the local computer.

   ```
   sftp> get picture.png /users/annejohnson picture.png
   ```
6. To disconnect and end the SFTP session, enter `exit` at the prompt.

Securing VPN Service

Virtual Private Network (VPN) is two or more computers or networks (nodes) connected by a private encrypted secure tunnel. This link simulates a local connection, as if the remote computer were attached to the local area network (LAN).

VPNs stress security by strong authentication of identity and encrypted data transport between the nodes for data privacy and inalterability. You'll be able to enable either or both of the encrypted transport protocols. Each has its own strengths and requirements.

Enabling Layer Two Tunneling Protocol, Secure Internet Protocol

Layer Two Tunneling Protocol over Secure Internet Protocol (L2TP/IPSec), uses strong IPSec encryption to "tunnel" data to and from the network nodes. IPSec requires security certificates (from a certificate authority like VeriSign), or a predefined shared secret between connecting nodes. A self signed certificate cannot be used. You can be your own certificate authority (using Certificate Assistant) and create your own root, server and client certificates. The server certificate for IPSec must contain the server's IP address and DNS name in the subject alternate name field and must be installed in the system keychain of the server. Client certificates must be installed in the system keychain of each client. The root certificate (your own or from a third party) which was used to sign your server and client certificates must be installed in the X509Anchors keychain on the server and client computers.

If a shared secret is used, it must be entered on the server as well as all clients. The shared secret is used by IPSec key management to authenticate the peers and as a factor in the generation of the encryption keys.

L2TP is the preferred VPN protocol due to its superior transport encryption and its ability to be authenticated using Kerberos.

Use Server Admin to designate L2TP as the transport protocol. By enabling this protocol, you must also configure the connection settings. You must designate an IPSec shared secret (if you don't use a signed security certificate), the IP address allocation range to be given to your clients, and group to be allowed VPN privileges (if desired). If both L2TP and PPTP are used, each protocol should have a separate, nonoverlapping address range.

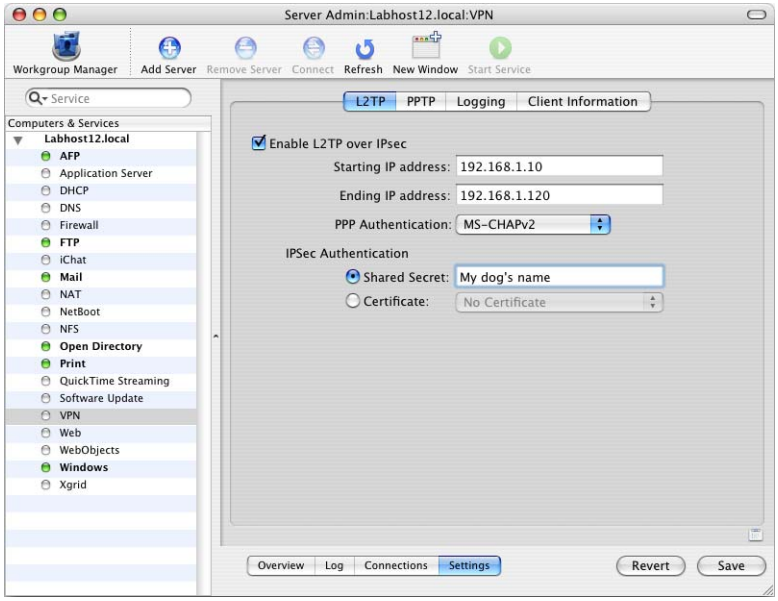

To enable L2TP:

1. In Server Admin, select the VPN Service in the Computers & Services list.
2. Click Settings.
3. Click L2TP and select "Enable L2TP over IPSec."
4. Set the beginning IP address of the allocation range.
5. Set the ending IP address of the allocation range.
6. Choose a PPP Authentication type.

   If your computer is bound to a Kerberos authentication server, choose Kerberos, otherwise choose MS-CHAPv2.
7. Either select Shared Secret and enter a shared secret, or select Certificate and choose a certificate.

   A shared secret is a string of text that the VPN service expects before it will receive the user name and password. A signed secure certificate is more secure than a shared secret.
8. Click Save.

Enabling and Configuring Point-to-Point Tunneling Protocol

Point-to-Point Tunneling Protocol (PPTP) is a common VPN protocol as well as the Windows standard VPN protocol. PPTP offers good encryption (provided the passwords used are strong passwords) and supports a number of authentication schemes. It uses the user-provided password to produce an encryption key. You can also allow 40-bit (weak) security encryption in addition to the default 128-bit (stronger) encryption if needed by your VPN clients. PPTP is necessary if you have old Windows clients or Mac OS X 10.2.x clients.

Use Server Admin to designate PPTP as the transport protocol. By enabling this protocol, you must also configure the connection settings. You should designate an encryption key length (40-bit in addition to 128-bit), the IP address allocation range to be given to your clients, and group to be allowed VPN privileges (if desired). If both L2TP and PPTP are used, each protocol should have a separate, nonoverlapping address range.

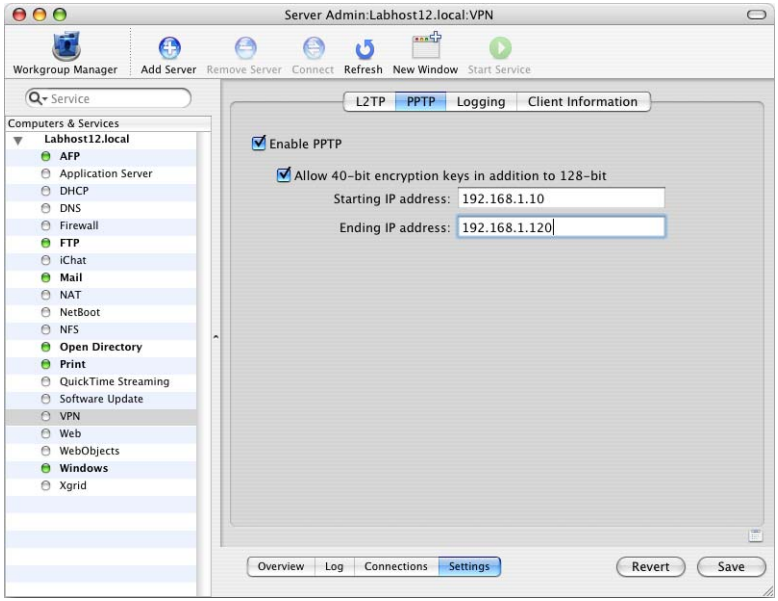

To enable PPTP:

1. In Server Admin, select the VPN Service in the Computers & Services list.
2. Click Settings.
3. Click PPTP.
4. Select "Enable PPTP."
5. Deselect "Allow 40-bit encryption keys in addition to 128-bit."

   Some VPN client applications require that you allow 40-bit encryption keys

6 Set the starting and ending IP addresses of the allocation range.

7 Click Save.

Authentication Methods

Mac OS X Server L2TP VPN uses either Kerberos v5 or Microsoft's Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) for authentication. Mac OS X Server PPTP VPN uses MS-CHAPv2, exclusively, for authentication.

Kerberos is a secure authentication protocol which depends on a Kerberos Key Distribution Server as a "trusted third party" to authenticate a client to a server. MSCHAPv2 authentication doesn't require the same authentication infrastructure as Kerberos. It encodes passwords when they're sent over the network and stores them in a scrambled form on the server offering good security during network transmission. It is also the standard Windows authentication scheme for VPN.

Both L2TP and PPTP VPN can use additional authentication methods. Each has its own strengths and requirements. It is not possible to choose any other authentication method for L2TP and PPTP using Server Admin. If you must configure an authentication scheme other than the default (for example, to use the RSA Security SecurID authentication), you need to edit the VPN configuration file manually. The configuration file is /Library/Preferences/SystemConfiguration/com.apple.RemoteAccessServers.plist.

Offering SecurID Authentication with VPN Service

RSA Security provides strong authentication through their product offering. They use hardware and software tokens to verify user identity. SecurID authentication is available for both L2TP and PPTP transports. For details and product offerings, see www.rsasecurity.com.

Mac OS X Server VPN service can offer SecurID authentication, but it cannot be set up from within Server Admin. You can use Server Admin to configure standard VPN services, but Server Admin does not have an interface for choosing your authentication method. If you must designate an authentication scheme (such as RSA Security SecurID) other than the default, you need to change the VPN configuration manually.

For additional information, see the *RSA SecurID Ready Implementation Guide,* located on the web at rsasecurity.agora.com/rsasecured/guides/imp_pdfs/MacOSX_ACE_51.pdf.

To manually configure RSA Security SecurID authentication:

1. Open Terminal.
2. Create a folder named /var/ace on your Mac OS X Server.

   ```
   $ sudo mkdir /var/ace
   ```

   Authenticate, if requested.
3. In Finder, choose Go > Go to Folder.
4. Type `/var/ace`.
5. Click Go.
6. Copy the sdconf.rec file from a SecurID server to /var/ace/.

   You will see a dialog indicating that the /var/ace/ folder cannot be modified. Click Authenticate to allow the copy.
7. Configure the VPN service (PPTP or L2TP) on your Mac OS X Server to enable EAP-SecurID authentication for the protocols you want to use it with.

   Enter the following in Terminal, replacing *protocol* with either `pptp` or `l2tp`:

   ```
   $ sudo serveradmin settings
       vpn:Servers:com.apple.ppp.protocol:PPP:AuthenticatorEAPPlugins:_array_i
       ndex : 0 = "EAP-RSA"
   $ sudo serveradmin settings
       vpn:Servers:com.apple.ppp.protocol:PPP:AuthenticatorProtocol:_array_ind
       ex: = "EAP"
   ```

   The remainder of Mac OS X Server VPN service configuration can be done using the Server Admin application.

Configuring Access Warning Banners

You can use a login window banner or a Terminal warning to provide notice of a computer's ownership to warn against unauthorized access or to remind authorized users of their consent to monitoring. For more information, see "Configuring Access Warnings" on page 59.

Securing Apple Remote Desktop

Apple Remote Desktop is an easy-to-use, powerful, open standards-based, desktop management tool. It provides several security mechanisms, which include AES-128 encryption that ensure you can use the tool securely and that all data is securely transferred to and from client and administrator computers.

For more information, see the Apple Remote Desktop administration guide.

Apple Remote Desktop is enabled by default in Sharing preferences. This service should be disabled unless your organization requires it. For more information, see "Securing Sharing Preferences" on page 105.

Encrypting Observe and Control Network Data

Although Remote Desktop sends authentication information, keystrokes, and management commands encrypted by default, you may want additional security. You can choose to encrypt all Observe and Control traffic, at a certain performance cost.

Encryption is done using an SSH tunnel between the participating computers. In order to use encryption for Observe and Control tasks, the target computers must have SSH enabled ("Remote Login" in the computer's Sharing Preference pane). Additionally, firewalls between the participating computers must be configured to pass traffic on TCP port 22 (SSH well known port).

If the you are trying to control a VNC server which is not Remote Desktop, it will not support Remote Desktop keystroke encryption. If you try to control that VNC server, you will get a warning that the keystrokes aren't encrypted which you will have to acknowledge before you can control the VNC server. If you chose to encrypt all network data, then you will not be able to control the VNC server because Remote Desktop is not able to open the necessary SSH tunnel to the VNC server.

To enable Observe and Control transport encryption:

1. Choose Remote Desktop > Preferences.
2. Click the Security button.
3. In the "Controlling computers" section, select "Encrypt all network data."

Encrypting Network Data During File Copy and Package Installations

Remote Desktop can send files for Copy Items and Install Packages via encrypted transport. This option is not enabled by default, and you must either enable it explicitly for each copy task, or in a global setting in Remote Desktop's preferences. Even installer package files can be intercepted if not encrypted.

To encrypt individual file copying and package installation tasks:

- In the Copy Items task or Install Packages task configuration window, select "Encrypt network data."

To set a default encryption preference for file copies:

1. In the Remote Desktop Preferences window, select the Security pane.
2. Check "Encrypt transfers when using Copy Items," or "Encrypt transfers when using Install Packages" as desired.

   Alternatively, you could encrypt a file archive before copying it. The encrypted archive could be intercepted, but it would be unreadable.

Securing Remote Apple Events

If you enable Remote Apple Events, you are allowing your computer to respond to events sent by other computers on your network. These events include AppleScript programs. A malicious AppleScript program can do things like delete your ~/Documents/ folder.

Remote Apple Events are disabled by default in Sharing preferences. This service should remain disabled unless your organization requires it. If you must enable this service, make sure you are on a trusted private network and disable it immediately after disconnecting from the network. For more information, see "Securing Sharing Preferences" on page 105.

Securing Network and Host Access Services

11

You can tailor Network and host access services in Mac OS X Server to protect your computer and network users. Proper configuration of services is important and will help create a hardened shell protecting your network.

Mac OS X Server includes several network and host access services that help you manage and maintain your network. This section describes recommended configurations for securing your network services.

For additional information about configuring network and host access services, see the network services administration guide.

Using IPv6 Protocol

IPv6 is short for Internet Protocol version 6. IPv6 is the Internet's next-generation protocol designed to replace the current Internet Protocol, IP Version 4 (IPv4 or just IP).

IPv4 has a limited number of possible unique addresses and is becoming obsolete as the Internet continues grow. IPv6 fixes this address problem and improves routing and network autoconfiguration. The increased number of network addresses eliminates the need for NAT. IPv6 is expected to gradually replace IPv4 over a number of years, with the two coexisting during the transition.

Mac OS X Server fully supports IPv6, which is configurable from Network preferences. You should disable the IPv6 protocol if your server and clients do not require it. Disabling the protocol prevents potential vulnerabilities on your computer. For information about disabling IPv6, see "Securing Network Preferences" on page 100.

To enable IPv6:

1. Open Network preferences.
2. In the Show pop-up menu, choose Built-in Ethernet.
3. Click TCP/IP.
4. Click Configure IPv6.

5 In "Configure IPv6," choose Automatically.

If you choose Manually, you will need to know your assigned IPv6 address, your router's IP address, and a prefix length.

6 Click OK.

7 Click Apply Now.

IPv6-Enabled Services

The following services in Mac OS X Server currently support IPv6 addressing:

- DNS (BIND)
- IP Firewall
- Mail (POP/IMAP/SMTP)
- SMB/CIFS
- Web (Apache 2)

Additionally, a number of command-line tools are installed with Mac OS X Server that support IPv6 (for example, `ping6` and `traceroute6`).

Note: The services listed above don't support IPv6 addresses in the user interface. They can be configured with command-line tools to add IPv6 addresses, but those same addresses fail if entered into address fields in Server Admin.

For more information about Internet Protocol version 6, see www.ipv6.org.

Securing DHCP Service

Mac OS X Server includes dynamic host configuration protocol (DHCP) service software, which allows it to provide IP addresses, LDAP server information, and DNS server information to clients.

Disabling Unnecessary DHCP Services

Using DHCP is not recommended. Assigning static IP addresses eases accountability and mitigates the risks posed by a rogue DHCP server. If DHCP use is necessary, only one system should act as the DHCP server and the service should be disabled on all other systems.

To disable the DHCP service:

1 Open Server Admin.

2 Select DHCP in the list for the server you're configuring.

3 Verify that the top of the window says "DHCP Service is: Stopped." If not, click "Stop Service."

Configuring DHCP Services

If you need to use the system as a DHCP server, don't distribute DNS, LDAP, and WINS information. To prevent distributing this information as part of DHCP, configure the DHCP service.

To configure the DHCP service:

1. Open Server Admin.
2. Select DHCP in the list for the server you're configuring.
3. Click Settings.
4. In the list that appears, select the subnet and click Edit.
5. Click DNS.
6. Delete any Name Servers listed.
7. Click LDAP.
8. Delete any server information that appears.
9. Click WINS.
10. Delete the WINS information.
11. Click the back arrow at the top left, and repeat steps 4 through 10 for any other subnets.
12. Click Save.

Assigning Static IP Addresses Using DHCP

You can assign a static address to a computer, if desired. This allows you to keep the ease of configuration of using DHCP, while allowing to you to have some static servers or services.

A static map consist of a specific IP address assigned to a network device. By using static maps, you avoid potential address conflicts and prevent hackers from easily obtaining valid IP addresses. With static maps, administrators can easily track network activity.

To assign a static IP address to a device, you need the device's Ethernet Address (sometimes called its MAC address or hardware address). Each network interface has its own Ethernet address.

Be aware that if you have a computer that moves from being wired to the network to a wireless network, it uses two different Ethernet addresses, one for the wired connection, and one for the wireless connection.

To assign a static IP address:

1. In Server Admin, select DHCP in the Computers & Services list.
2. Click Settings.

3 Click Static Maps.

4 Click the Add (+) button.

To modify an existing static map, click the Edit (/) or the Delete (–) button.

5 Enter the Ethernet Address of the computer that is to get a static address.

6 Enter the IP address you want to assign to it.

7 Enter the name of the computer.

8 Click OK.

9 Click Save.

Securing DNS Service

Mac OS X Server includes an installation of Berkeley Internet Name Daemon (BIND) 9.2 for use as domain name server software. BIND is an open source implementation and is used by the majority of name servers on the Internet. The DNS server software should be deactivated if your server is not intended to be a DNS server.

To disable the DNS service:

1 Open Server Admin.

2 Click DNS in the list for the server you're configuring.

3 Verify that the top of the window says "DNS Service is: Stopped." If not, click "Stop Service."

Understanding BIND

BIND implements DNS through the *name daemon* or /etc/named.conf file. Some configurations require you to modify the BIND configuration files. You should have a thorough understanding of DNS before you attempt to modify BIND configuration files.

If you edit named.conf to configure BIND, make sure that you don't change the controls statement `inet` settings. Otherwise, Server Admin will be unable to retrieve status information for DNS.

The `inet` settings should look like this:

```
controls {
     inet 127.0.0.1 port 54 allow {any;}
     keys { "rndc-key"; };
};
```

Also, note that using Server Admin after editing the BIND configuration files might overwrite some changes.

For more information about DNS and BIND, see the following:

- *DNS and BIND, 4th edition*, by Paul Albitz and Cricket Liu (O'Reilly and Associates, 2001)
- The International Software Consortium website: www.isc.org and www.isc.org/products/BIND/
- The DNS Resources Directory: www.dns.net/dnsrd/

Turning Off Zone Transfers and Recursive DNS Queries

If DNS use is necessary, only one server should act as the DNS server and the service should be disabled on all other servers. Use Server Admin to configure a DNS server. Unless your site requires them, turn off Zone Transfers and recursive DNS queries.

To turn off zone transfers and recursive DNS queries:

1. Open Server Admin.
2. Click DNS in the list for the server you're configuring.
3. Click Settings.
4. Deselect "Zone transfers."

   Zone transfers, if needed, should be set up so that they only occur between trusted servers. This requires manually editing the BIND configuration files.
5. Deselect "Recursion."

   If your site requires recursion, you should allow recursive queries only from trusted clients and not from any external networks.
6. Click Save.

Make sure that both forward and reverse zones are established and fully populated. Otherwise, any Open Directory server using the DNS service will not work correctly.

Disabling Recursion

Recursion is the process of fully resolving domain names into IP addresses. Users' applications depend on the DNS server to perform this function. Other DNS servers that query yours don't have to perform the recursion.

To prevent malicious users from altering the primary zone's records ("cache poisoning") or allowing unauthorized use of the server for DNS service, you can disable recursion. However, if you disable it, your own users are not able to use your DNS service to look up any names outside of your zones.

You should disable recursion only if no clients are using this DNS server for name resolution and no servers are using it for forwarding.

To disable recursion:

1. In Server Admin, select DNS in the Computer & Services list.
2. Click Settings.
3. Select General.
4. Deselect Recursion to disable it.

If you choose to enable recursion, consider disabling it for external IP addresses, but enabling it for internal IP addresses. This requires manually editing the BIND configuration files.

Understanding DNS Security

DNS servers are susceptible to several kinds or attacks. By taking extra precautions, you can prevent the problems and downtime associated with malicious users. Several kinds of security hacks are associated with DNS service:

- DNS spoofing
- Server mining
- DNS service profiling
- Denial of Service (DoS)
- Service piggybacking

DNS Spoofing

DNS spoofing is adding false data into the cache of the DNS server. This allows hackers to do any of the following:

- Redirect real domain name queries to alternative IP addresses. For example, a falsified address for a bank could point a computer user's browser to a different IP address that is controlled by the hacker. A duplicate website could fool the user into giving his or her bank account numbers and passwords to the hacker.

  Also, a falsified mail address could allow a hacker to intercept mail sent to or from a domain. If the hacker also forwards those emails to the correct mail server after copying them, this can go undetected indefinitely.
- Prevent proper domain name resolution and access to the Internet. This is the most benign of DNS spoof attacks. It merely makes a DNS server appear to be malfunctioning.

The most effective method to guard against these attacks is vigilance. This includes maintaining up-to-date software, as well as auditing your DNS records regularly. As exploits are found in the current version of BIND, the exploit is patched and a security update is made available for Mac OS X Server. Apply all such security patches. Regular audits of your DNS records can help prevent these attacks.

Server Mining

Server mining is the practice of getting a copy of a complete primary zone by requesting a zone transfer. In this case, a hacker pretends to be a secondary zone to another primary zone and requests a copy of all of the primary zone's records.

With a copy of your primary zone, the hacker can see what kinds of services a domain offers, and the IP address of the servers that offer them. The hacker can then try specific attacks based on those services.

To defend against this attack, you must specify which IP addresses are allowed to request zone transfers (your secondary zone servers) and disallow all others. Zone transfers are accomplished over TCP on port 53. Limiting zone transfers involves blocking zone transfer requests from anyone but your secondary DNS servers.

To specify zone transfer IP addresses, create a firewall filter that allows only IP addresses inside your firewall to access TCP port 53. Follow the instructions in "Creating Advanced Firewall Rules" on page 217, using the following settings:

- Allow packet
- Port 53
- TCP protocol
- Source IP is the IP address of your secondary DNS server
- Destination IP is the IP address of your primary DNS server

DNS Service Profiling

Another common reconnaissance technique used by malicious users is to profile your DNS service. First a hacker makes a BIND version request. The server reports what version of BIND is running. The hacker then compares the response to known exploits and vulnerabilities for that version of BIND.

To defend against this type of attack, you can configure BIND to respond with something other than the current version.

To alter BIND's version response:

1. Open the /etc/named.conf file from Terminal.

   ```
   $ sudo pico /etc/named.conf
   ```

2. Add the following to the "options" brackets of the configuration file.

   ```
   version "[enter your text here]";
   ```

   where `enter your text here` could be something like `Hidden Information`.

3. Save the configuration file.

Denial of Service

The Denial of service (DoS) attacks are very common and easy to do. A hacker sends many service requests and queries, overloading the server, forcing the server to use all of its processing power and network bandwidth trying to respond. This attack prevents legitimate use of the service by overloading it.

It is difficult to prevent this type of attack before it begins. Constant monitoring of the DNS service and server load allows an administrator to catch the attack early and mitigate its damaging effect.

The easiest way to guard against this type of attack is to block the offending IP address with your firewall. See "Creating Advanced Firewall Rules" on page 217. Unfortunately, this means the attack is already underway, the hacker's queries are being answered, and the activity logged.

ARP spoofing

This type of attack, also known as ARP poisoning, allows an attacker to take over a computers IP address by modifying its ARP caches. The attacker must be on either the same network as the computer it is attacking or host that the computer is communicating with. The attacker can also use ARP spoofing for a man-in-the-middle attack, which forwards all traffic from a computer to the attacker's computer. This allows the attacker to view packets and look for passwords and confidential data. ARP spoofing can also be used to create a DoS attack stopping all network traffic.

By configuring your network with static IP addresses and monitoring your network traffic you can keep unauthorized users from maliciously using your network.

Service Piggybacking

This type of attack is not typically done by malicious intruders, but by common Internet users who learn this trick from other users. They might feel that their DNS response time with their own Internet service provider is too slow. The Internet users configure their computer to query another DNS server instead of their own ISP's DNS servers. Effectively, more users are accessing the DNS server than have been planned for.

You can guard against this by limiting or disabling DNS recursion. If you plan to offer DNS service to your own LAN users, they need recursion to resolve domain names, but you don't want to provide this service to any other Internet users.

To prevent recursion entirely, see "Disabling Recursion" on page 209.

The most common balance is allowing recursion for requests coming from IP addresses within your own range, but denying recursion to external addresses. BIND allows you to specify this in its configuration file, named.conf. Edit your named.conf file to include the following:

```
options {
...
  allow-recursion{
     127.0.0.0/8;
     [your internal IP range of addresses, like 192.168.1.0/27];
               };
     };
```

Securing Firewall Service

Firewall service is software that protects the network applications running on Mac OS X Server. Turning on firewall service is similar to erecting a wall to limit access. Firewall service scans incoming IP packets and rejects or accepts these packets based on the set of rules you create. You can restrict access to any IP service running on the server, and you can customize rules for all incoming clients or for a range of client IP addresses.

Important: The firewall service can disrupt network communications and its configuration can be tricky to implement. Do not implement recommendations without understanding their purpose or impact.

Services, such as web and FTP are identified on your server by a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port number. When a computer tries to connect to a service, firewall service scans the rule list for a matching port number.

The default firewall configuration on Mac OS X Server denies access to all but a few TCP services and blocks access to UDP services, except for responses to DNS queries. The goal of configuring the firewall is to identify and permit only those hosts and services you would like to allow, and then deny all others. The recommended settings deny all TCP and UDP services except those explicitly allowed.

Important: You should not perform any server configuration remotely—particularly the firewall service, because of the risk of disabling communications to the remote host.

Planning Firewall Setup

Plan your IP firewall service by deciding which services you want to provide access to. Mail, web, and FTP services generally require access by computers on the Internet. File and print services are most likely restricted to your local subnet.

Once you decide which services you want to protect using firewall service, you must determine which IP addresses you want to allow access to your server and which IP addresses you want to deny access to your server. You can then create the appropriate rules.

Once the Firewall service is configured, network users may request that the rules be changed to allow additional services. These changes should be resisted and an approval process be put in place to monitor these changes.

Starting the Firewall Service

In Server Admin, select Firewall and click Start Service. By default, this blocks all incoming ports except those used to configure the server remotely. If you're configuring the server locally, turn off external access immediately.

Important: If you add or change a rule after starting firewall service, the new rule affects connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server are disconnected.

Although the firewall is treated as a service by the Server Admin application, it is not implemented by a running process like other services. It is simply a set of behaviors in the kernel, controlled by the `ipfw` and `sysctl` tools.

To start and stop the firewall, the Server Admin application sets a switch using the `sysctl` tool. When the computer starts, a startup item named IPFilter checks the /etc/hostconfig file for the "IPFILTER" flag. If it is set, the `sysctl` tool is used to enable the firewall:

```
$ sysctl -w net.inet.ip.fw.enable=1
```

Otherwise, it disables the firewall:

```
$ sysctl -w net.inet.ip.fw.enable=0
```

Note: The rules loaded in the firewall remain there regardless of this setting. It's just that they are ignored when the firewall is disabled.

Like most startup items, the IPFilter startup item launches in a predetermined order, and only after certain prerequisite startup items have completed. The login window is presented while startup items might be running. It is, therefore, possible to log in while the firewall has not been set to its configured settings. The startup item that sets up the firewall should finish within a few minutes of starting the system.

Creating an IP Address Group

By grouping IP addresses you are able to set Firewall rules for large numbers of network devices at a time and allow for much better organization. This enhances the security if your network.

These groups are used to organize and target the rules. The "any" address group is for all addresses. Two other IP address groups are present by default, intended for the entire "10.0.0.0" range of private addresses, and the entire "192.168.0.0" range of private addresses.

Addresses can be listed as individual addresses (192.168.2.2), IP address and CIDR notation (192.168.2.0/24), or IP address and netmask notation (192.168.2.0:255.255.255.0).

By default, an IP address group is created for all incoming IP addresses. Rules applied to this group affect all incoming network traffic.

To create an address group:

1. In Server Admin, select Firewall in the Computers & Services list.
2. Click Settings.
3. Select General.
4. Click the Add (+) button to the right the Address Group pane.
5. Enter a group name.
6. Enter the addresses and subnet mask you want the rules to effect.

   Use the Add (+), Edit (/), and Delete (–) buttons.

   Use the word "any" to indicate any IP address.
7. Click OK.
8. Click Save.

Creating Firewall Service Rules

A Firewall rule specifies which services to let through your firewall, and which ones to keep out. The rule defines the parameters that the Firewall compares against each connection, then determines what action to take for each connection.

You can activate rules based on address groups as destination IP numbers. By default, firewall service blocks incoming TCP connections on ports that are not essential for remote administration of the server and allows all UDP connections. Also by default, rules are in place that allow specific responses to outgoing requests. Before you turn on firewall service, make sure you've set up rules allowing access from IP addresses you choose; otherwise, no one will have access to your server.

You can easily allow standard services through the firewall without advanced and extensive configuration. Standard services include (but are not limited to):

- SSH access
- Web service
- Apple file service
- Windows file service
- FTP service
- Printer sharing
- DNS/Multicast DNS
- ICMP Echo Reply (incoming pings)
- VPN
- QTSS media streaming

Important: If you add or change a rule after starting firewall service, the new rule affects connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server are disconnected.

To open the firewall for standard services:

1. In Server Admin, select Firewall in the Computers & Services list.
2. Click Settings.
3. Select Services.
4. Select an address group from the Edit Services pop-up menu.
5. Choose either to allow all traffic for the address group or to allow traffic on designated points.
6. Check Allow for each service you want the address group to have access to. If you don't see the service you need, you can add a port and description to the services list.
7. Click Save.

Creating Advanced Firewall Rules

You use advanced rules to further configure all other services, strengthen your network security, and fine-tune your network traffic through the firewall. By default, all UDP are blocked, except those in response to an outgoing query. You should apply rules to UDP ports sparingly, if at all, because denying certain UDP responses could inhibit normal networking operations.

If you apple rules to UDP ports, don't select the "Log all allowed packets" option in the rule configuration windows in Server Admin. Since UDP is a "connectionless" protocol, every packet to a UDP port is logged if you select this option

You can use the Advanced Settings pane to configure specific rules for IP firewall. IP firewall rules contain originating and destination IP addresses with subnet masks. They also specify what to do with the network traffic received. You can apply a rule to all IP addresses, a specific IP address, or a range of IP addresses.

Addresses can be listed as individual addresses (192.168.2.2) or as ranges defined by an IP address and CIDR netmask (192.168.2.0/24).

To create an advanced IP firewall rule:

1. In Server Admin, what Firewall in the Computers & Services list.
2. Click Settings.
3. Click Advanced Rules.
4. Click the Add (+) button.

   Alternatively, you can select a rule similar to the one you want to create, click Duplicate and then click Edit.
5. Select whether this rule will allow or deny access in the Action pop-up menu.

   If you choose Other, enter the action desired (for example, log).
6. Choose a protocol from the Protocol pop-up menu.

   If you choose Other, enter the desired protocol (for example, icmp, esp, ipencap).
7. Choose a service from the pop-up menu.

   If you want to select a nonstandard service port, choose Other.
8. If desired, choose to log packets that match the rule.
9. Choose an address group from the pop-up menu as the source of filtered traffic.

   If you don't want to use an existing address group, enter the source IP address range (with CIDR notation) you want to filter.

   If you want the rule to apply to any address, choose "any" from the pop-up menu.
10. If you have selected a nonstandard service port, enter the source port number.

11 Choose an address group from the pop-up menu as the destination of filtered traffic.

If you don't want to use an existing address group, enter the destination IP address range (with CIDR notation).

If you want the rule to apply to any address, choose "any" from the pop-up menu.

12 If you have selected a nonstandard service port, enter the destination port number.

13 Choose which network interface this rule applies to.

"In" refers to the designated WAN interface.

"Out" refers to the designated LAN interface.

If you select Other, enter the interface name (en0, en1, fw1, and so on).

14 Click OK.

15 Click Save to apply the rule immediately.

Enabling Stealth Mode

You can hide the existence of your firewall by choosing not to send a connection failure notification to any connection that is blocked by the firewall. This effectively hides your server's closed ports. For example, if a network intruder tries to connect to your server, even if the port is blocked, he or she knows that there is a server and might find other ways to intrude. If stealth mode is enabled, instead of being rejected, he or she won't receive any indication that an attempted connection ever took place.

To enable stealth mode:

1 In Server Admin, select Firewall in the Computers & Services list.

2 Click Settings.

3 Select Advanced Rules.

4 Select "Enable for TCP" and/or "Enable for UDP," depending on your firewall configuration.

5 Click Save.

Setting Up Firewall Service Logging

Firewall service logging is important to security. With logs, administrators can monitor and track connections and attempted connections to the firewall.

You can log only the packets that are denied by the rules you set, only the packets that are allowed, or both. Both logging options can generate a lot of log entries, but there are ways to limit the volume:

- Log packets only as long as necessary.
- Limit the total number of packets, using the Logging Settings panel.
- Add a "count" rule in the Advanced settings to tally the number of packets that match the characteristics you're interested in measuring (such as attempted login to a specific port).

You can choose to log allowed packets, denied packets, and a designated number of packets.

To set up logs:

1. In Server Admin, select Firewall in the Computers & Services list.
2. Click Settings.
3. Select General.
4. Select the logging options you want.
5. Click Save to start logging.

Each rule you create in Server Admin corresponds to one or more rules in the underlying firewall software. Log entries show you the rule applied, the IP address of the client and server, and other information.

The log view shows the contents of /var/log/ipfw.log. You can filter the rules with the text filter box.

Securing NAT Service

Network Address Translation (NAT) is sometimes referred to as IP masquerading. NAT is used to allow multiple computers access to the Internet with a single assigned public or external IP address. This creates a more secure network environment by hiding the internal private IP addresses.

The NAT service further enhances security by limiting communication between your private network and a public network (such as the Internet):

- Communication from a computer on your private network is translated from a private IP address to a shared public IP address. Multiple private IP addresses are configured to use a single public IP address.
- Communication to your private network is translated and forwarded to an internal private IP address (IP forwarding). The external computer cannot determine the private IP address. This created a barrier between your private network and the public network.
- Communication from a public network cannot come into your private network unless it was requested. It is only allowed in response to internal communication.

The NAT service takes all the traffic from your private network and remembers which internal address made the request. When NAT receives the response to the request, it forwards it to the originating computers. Traffic that originates from the Internet does not reach any of the computers unless port forwarding is enabled.

Important: Firewall service must be enabled for NAT to function.

The NAT server software should be deactivated if your server is not intended to be a NAT server.

To disable the NAT service:

1. Open Server Admin.
2. Click NAT in the list for the server you're configuring.
3. Verify that the top of the window says "NAT Service is: Stopped." If not, click "Stop Service."

Configuring NAT Service

Use Server Admin to indicate which network interface is connected to the Internet or other external network.

To configure NAT service:

1. In Server Admin, select NAT in the Computers & Services list.
2. Click Settings.
3. Select "IP Forwarding and Network Address Translation."

4 Choose the network interface from the "Network connection to share" pop-up menu.

This interface should be the one that connects to the Internet or external network.

5 Click Save.

Configuring Port Forwarding

You can direct incoming traffic on your private network to a specific IP address. This allows you to set up computers on your private network that handle certain incoming connections without exposing the other computers to outside connections. For example, you could set up a web server behind NAT server and forward all incoming TCP connection requests on port 80 to the designated web server.

You can't forward the same port to multiple computers, but you can forward any number of different ports to the same computer.

Enabling port forwarding requires use of Terminal, as well as administrator access to root privileges through `sudo`. You must edit a plist, and the contents of that plist are used to generate the file /etc/nat/natd.conf.apple, which is passed to the NAT daemon when it is started. Do not try to edit /etc/nat/natd.conf.apple directly. If you choose to use a plist editor instead of a command-line text editor, you might need to alter the following instructions accordingly.

To forward port traffic:

1 If the file /etc/natd.plist doesn't exist, make a copy of the default NAT daemon plist.

```
$ sudo cp /etc/nat/natd.plist.default /etc/natd.plist
```

2 Using a command-line editor, add the following block of XML text to /etc/natd.plist before the last two line of the file (`</dict>` and `</plist>`). This will configure NAT to use the port mapping rule `redirect_port tcp 1.2.3.4:80 80.`

```
<key>redirect_port</key>
<array>
<dict>
<key>proto</key>
<string>TCP</string>
<key>targetIP</key>
<string>1.2.3.4</string>
<key>targetPortRange</key>
<string>80</string>
<key>aliasPortRange</key>
<string>80</string>
</dict>
</array>
```

3 Save your changes.

The changes made, except for those settings that Server Admin can change and comments, are respected by the server configuration tools (Server Admin, Gateway Setup Assistant, and serveradmin).

4 Confirm those settings using the serveradmin tool:

```
$ sudo serveradmin settings nat
...
nat:redirect_port:_array_index:0:proto = "tcp"
nat:redirect_port:_array_index:0:targetPortRange = "80"
nat:redirect_port:_array_index:0:aliasPortRange = "80"
nat:redirect_port:_array_index:0:targetIP = "1.2.3.4"
```

5 Configure NAT service in Server Admin. For more information, see "Configuring NAT Service" on page 220.

6 Click Save.

Securing Bonjour Service

Bonjour is a protocol for discovering file, print, chat, music sharing, and other services on IP networks. Bonjour listens for service inquiries from other computers, and also provides information regarding your available services.

Users and applications on your local network can use Bonjour to quickly determine which services are available on your computer. Although this might seem like a security risk, malicious intruders can use their own tools, such as port scanners, to locate the same services advertised by Bonjour. You should disable any unused services that you don't want others to discover through Bonjour.

Enter the following command to disable Bonjour:

```
$ sudo launchctl unload -w /System/Library/LaunchDaemons/
     com.apple.mDNSResponder.plist
```

You won't be able to use network printing using Bonjour, so you'll have to manually configure network printers. This can also disable some functionality in other applications that rely on Bonjour, or possibly make them unusable. For example, there are issues with calendar and address book sharing, and finding iChat buddies.

If disabling Bonjour causes vital applications to break, enter the following command to reenable Bonjour:

```
$ sudo launchctl load -w /System/Library/LaunchDaemons/
     com.apple.mDNSResponder.plist
```

If you decide to re-enable Bonjour, block UDP port 5353 on your firewall to block externally-originating Bonjour traffic.

Securing Collaboration Services

12

Collaboration services help users share information for increased productivity. Securing the access and transfer of shared information will protect your data.

Collaboration services promote interactions among users, facilitating teamwork and productivity. The iChat service provides a secure way for users to chat. In order to use iChat service on a particular server, users must be defined in directories the server uses to authenticate users. For more information about configuring search paths to directories, see the Open Directory administration guide.

For information about configuring collaboration services, see the collaboration services administration guide.

Disabling iChat Service

The iChat server software should be deactivated if your server is not intended to be an iChat server. Disabling the service prevents potential vulnerabilities on your computer.

To disable iChat service:

1. Open Server Admin.
2. Select the server you are configuring from the Computer & Services list.
3. Click iChat.
4. Click Overview and verify that the pane says "iChat Service is: Stopped." If not, click Stop.

Securely Configuring iChat Service

If your organization requires the use of iChat service, configure it to use SSL. SSL communication certifies the identity of the server and establishes secure, encrypted data exchange.

You identify an SSL certificate for iChat service to use the first time you set up iChat service, but you can use a different certificate later if you like. You can use a self-signed certificate or a certificate imported from a Certificate Authority. For more information about defining, obtaining, and installing certificates on your server, see "Readying Certificates" on page 174.

Sending messages to multiple recipients over an internal iChat sever does not require a .Mac identity. The internal iChat server (Jabberd) requires a server side SSL certificate which is utilized by each client to establish an SSL session (similar to a web access session). A .Mac certificate is required to establish encrypted sessions between two iChat clients communicating using text, audio, and video.

To securely configure iChat service:

1. Open Server Admin.
2. In the Computers & Services list, select iChat for the server you are configuring.
3. Click Settings.
4. The Host Domains field designates the domain names you want iChat service to support. Initially, the server's host name is displayed. You can add other names that resolve to the iChat server's IP address, such as aliases defined in DNS.

   Host domains are used to construct screen names, which identify iChat service users. An example of a screen name is annejohnson@example1.apple.com.
5. The text in the Welcome Message window is displayed by chat clients when they connect to iChat service. Change the default text displayed in the field if required.
6. From the SSL Certificate pop-up menu, choose a Secure Sockets Layer (SSL) certificate you want iChat service to use. The menu lists all SSL certificates that have been installed on the server.
7. Click Save, and then click Start Service.
8. Make sure the iChat server's Open Directory search path includes directories in which the users and group members that you want to communicate using iChat service are defined. The Open Directory administration guide explains how to set up search paths.

   Any user or group member defined in the Open Directory search path is now authorized to use iChat service on the server, unless you deny them access to iChat service.

For additional security enhancements, you can further restrict the iChat service by using SACLs and firewall rules. These are configured based on your organizations network environment.

- You can configure SACLs to restrict iChat access to specific users or groups. For more information about configuring SACLs, see "Setting Service Access Privileges" on page 188.
- You can configure firewall rules that prevent iChat connections from unintended sources. For more information, see "Creating Firewall Service Rules" on page 216.

Viewing iChat Service Logs

iChat service logging is important for security. With logs, you can monitor and track communication through the iChat service. The iChat service log, /var/log/system.log, can be accessed using Server Admin.

To view the iChat service log:

1. Open Server Admin.
2. In the Computers & Services list, select iChat for the server you are logging.
3. Click Log.

Securing Mail Service

13

Mail service is crucial in today's dispersed work environments. Protect your mail with encryption, adaptive junk mail filtering, and virus detection.

Mail service in Mac OS X Server allows network users to send and receive email over your network or across the Internet. This creates vulnerabilities for your server.

The Mail service in Mac OS X Server consists of three components, all based on open standards with full support for Internet mail protocols:

- Postfix, the SMTP mail transfer agent (outgoing mail)
- Cyrus, which supports IMAP and POP (incoming mail)
- Mailman, which provides mailing list management features

For more information about configuring mail service, see the mail service administration guide.

Disabling Mail Service

The mail service software should be disabled if your server is not intended to be a mail server. Disabling the service prevents potential vulnerabilities on your server. To disable the mail service, you must turn off support for the IMAP, SMTP, and POP protocols that are not required. The mail service is disabled by default, but verification is recommended.

To deactivate unnecessary mail protocols:

1. Open Server Admin.
2. In the Computers & Services list, select Mail for the server you are configuring.
3. Click Overview and verify that the pane says "Mail Service is: Stopped." If not, click Stop Service.
4. Click Settings.
5. Deselect "Enable SMTP" if the system will not be used as an outgoing mail server.

6 Deselect "Enable IMAP" if the system will not be used as an incoming mail server.

7 Deselect "Enable POP" if the system will not be used as an incoming mail server.

8 Click Save.

Configuring Mail Service for SSL

If any mail service protocols are required, their communications should be protected by SSL. Enabling SSL for incoming and outgoing mail service causes communications between the mail server and its clients to be encrypted, protecting clients from eavesdroppers on the local network. You should also use different servers for providing outgoing mail service and incoming mail service where possible.

The steps to create SSL certificates for the mail server are similar to those for the web server. If the mail service and web service exist on the same server and use the same domain name, the same server certificate could be used for both services. However, this is not recommended.

The following steps describe the command-line method for creating certificates. For information about defining, obtaining, and installing certificates on your server using Certificate Manager in Server Admin, see "Readying Certificates" on page 174.

To securely configure mail service with SSL:

1 Open Terminal and go to the /usr/share/certs/ folder.

If the /usr/share/certs folder does not exist, create it.

2 Create a key pair for the mail server using the `openssl` command.

```
$ sudo openssl genrsa -out mailserver.key 2048
```

This differs from the web server certificate in that it is not encrypted (no `-des3` option). The mail server requires an unencrypted key.

3 Create the CSR with the mail server key:

```
$ sudo openssl req -new -key mailserver.key -out mailserver.csr
```

4 Fill out the following fields as completely as possible:

```
Country Name:
State or Province Name:
Locality Name (city):
Organization Name:
Organizational Unit Name:
Common Name:
Email Address:
```

The Common Name field is critical. It must match the domain name of the mail server exactly, or the certificate will not work.

5 Sign the mailserver.csr certificate.

```
$ openssl ca -in mailserver.csr -out mailserver.crt
```

6 The mail server expects the key and certificate inside the same file, so concatenate the key and certificate:

```
$ sudo -s
$ cat mailserver.key mailserver.crt > mailserver.pem
$ exit
```

This creates the mailserver.pem file. This file can be moved to the mail server and installed.

7 Install a mail certificate for the outgoing mail service protocol.

If you're running an outgoing mail service protocol and have decided to act as your own CA, copy the mailserver.pem file to the /etc/postfix/ folder.

```
$ cp /usr/share/certs/mailserver.pem /etc/postfix/
```

8 Change the name to server.pem.

```
$ mv /etc/postfix/mailserver.pem /etc/postfix/server.pem
```

If you've purchased a certificate from a commercial CA, follow their instructions to ensure that the correct information ends up in /etc/postfix/server.pem.

9 Install a mail certificate for an incoming mail service protocol.

If you're running an incoming mail service protocol and have decided to act as your own CA, copy the mailserver.pem file to the /var/imap/ folder.

```
$ cp /usr/share/certs/mailserver.pem /var/imap/
```

10 Change the name to server.pem.

If you've purchased a certificate from a commercial CA, follow their instructions to ensure that the correct information ends up in /var/imap/server.pem.

```
$ mv /var/imap/mailserver.pem /var/imap/server.pem
```

11 Change the ownership of the server.pem file so the IMAP and POP server can read it:

```
$ chown cyrus /var/imap/server.pem
```

12 Open Server Admin.

13 In the Computers & Services list, select Mail for the server you are configuring.

14 Click Settings.

15 Click Advanced.

16 Click Security.

17 Under Secure Sockets Layer (SSL), choose "Require" from the SMTP SSL pop-up menu.

18 Under Secure Sockets Layer (SSL), choose "Require" from the IMAP and POP SSL pop-up menu.

19 Click Save.

Three options exist for the server's SSL support: Require, Use, and Don't Use. "Require" is the recommended option. "Use" allows both regular and SSL connections and is better than "Don't Use." Remember that SMTP mail clients must support SSL connections in addition to setting this up on the mail server. On a homogeneous Mac OS X Server network, this isn't an issue since the Mail client supports SSL, but on a heterogeneous network, SSL support on the client side might not exist.

Important: Mail clients must be set up to use SSL connections. Configuring an active mail server in the manner described causes a loss of service until the clients are reconfigured. Setting the "Use" option for a small period of time to allow clients to switch before "Require" is set might help them avoid a denial of service.

Configuring Authentication Support

Mail service protocols use authentication to protect users' passwords by requiring that connections use a secure method of authentication. When a user connects with secure authentication, the user's mail client software encrypt the users' password before sending it to your mail service.

Each mail service protocol has a set of supported methods of authentication. Authentication support protects users' mail passwords as they travel across the network. Although a proper SSL setup already encrypts the mail client-server communications, you should also use a secure authentication method.

Enabling authentication will:

- Make your users authenticate with their email client before accepting any mail to send.
- Frustrate mail server abusers trying to send mail through your system without your consent.

If you don't choose any method of SMTP authentication or don't authorize specific SMTP servers to relay for, the SMTP server allows anonymous SMTP mail relay, which is considered an "open relay." Junk mail senders can exploit the open relays to hide their identities and send illegal junk mail without penalty.

There is a difference between relaying mail and accepting delivery of mail:

- Relaying mail means passing mail from one (possibly external) mail server or a local user's email client to another (third) mail server.
- Accepting delivery means receiving mail from a (possibly external) mail server to be delivered to the server's own email users.

Mail addressed to local recipients is still accepted and delivered. Enabling authentication for SMTP requires authentication from any of the selected authentication methods prior to relaying mail. SMTP authentication is used in conjunction with restricted SMTP mail transfer to limit junk mail propagation.

If your computer is integrated into a Kerberos realm, select "Kerberos" for the mail service protocol (SMTP, IMAP, or POP) your system offers. Before enabling Kerberos authentication for incoming mail service, you must integrate with a Kerberos server. If you're using Mac OS X Server for Kerberos authentication, this is already done for you.

If your computer is not integrated into a Kerberos realm, select "CRAM-MD5" in the SMTP and IMAP columns, and "APOP" in the POP column. If you configure your mail service to require CRAM MD-5, mail users' accounts must be set to use a password server that has CRAM MD-5 enabled. For more information, see the Open Directory administration guide.

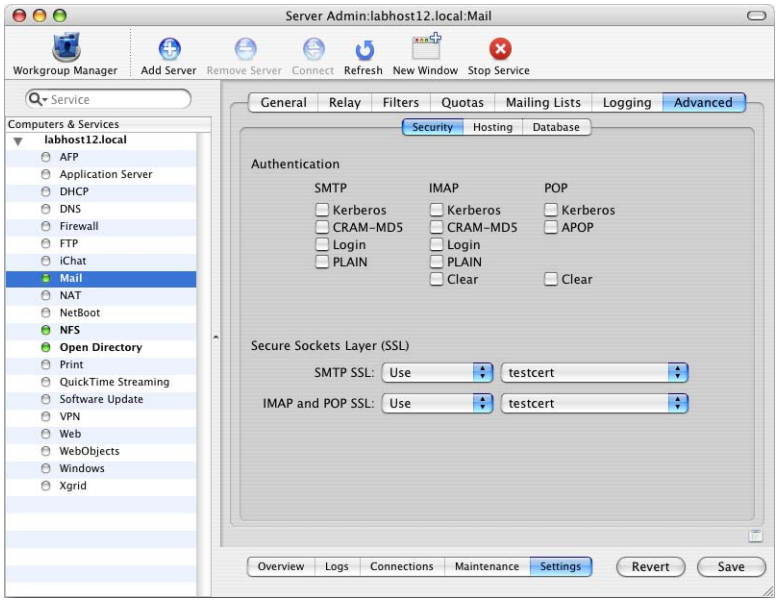

To set mail authentication:

1 Open Server Admin.

2 In the Computers & Services list, select Mail for the server you are configuring.

3 Click Settings.

4 Click Advanced.

5 Deselect all options in the Authentication section.

6 Select the required authentication method for each mail service protocol you are configuring.

Kerberos authentication should be used for each mail service. If your network does not have a Kerberos server available, use the challenge-response authentication mechanism message-digest algorithm 5 (CRAM-MD5), to secure authentication. You should not use PLAIN or Clear Authentication without an SSL certificate.

7 Click Save.

Restricting SMTP Relay

Your mail service can restrict SMTP relay by allowing only approved hosts to relay mail. You create the list of approved servers. Approved hosts can relay through your mail service without authenticating. Servers not on the list cannot relay mail through your mail service unless they authenticate first. All hosts, approved or not, can deliver mail to your local mail users without authenticating.

To restrict SMTP relay:

1. Open Server Admin.
2. In the Computers & Services list, select Mail for the server you are configuring.
3. Click Settings.
4. Click Relay.
5. Select "Accept SMTP relays only from these."
6. Edit the list of hosts.
7. Click the Add (+) button to add a host to the list. You can use a variety of notations.

   Enter a single IP address or the network/netmask pattern (such as 192.168.40.0/21).

   Enter a host name, such as mail.example.com.

   Enter an Internet domain name, such as example.com.
8. Click the Edit (/) or Delete (–) button to change or delete the currently selected host.
9. Click Save.

Enabling Mail Filtering

Once a mail delivery connection is made and the message is accepted for local delivery (relayed mail is not screened), the mail server can screen it before delivery. Mac OS X Server uses SpamAssassin to analyze the text of a message and gives it a probability rating for being junk mail.

No junk mail filter is 100% accurate in identifying unwanted email. It's for this reason that the junk mail filter in Mac OS X Server doesn't delete or remove junk mail from being delivered. Instead, it marks the mail as potential junk mail. The user can then decide if it's really unsolicited commercial email and deal with it accordingly. Many email clients even use the ratings that SpamAssassin adds as a guide in automatically classifying the mail for the user.

For more information about SpamAssassin, see spamassassin.apache.org.

To enable junk mail filtering:

1. Open Server Admin.
2. In the Computers & Services list, select Mail for the server you are configuring.
3. Click Settings.
4. Select Filters.
5. Select "Scan email for junk mail."
6. Set the "Minimum junk mail score:" to Least, Moderate, or Most.

   The junk mail score meter sets how many junk mail flags can be applied to a single message before it is processed as junk mail. If you set it to "Least," any mildly suspicious email is tagged and processed as junk mail. If you set it to "Most" it takes a high score (in other words, a lot of junk mail characteristics) to mark it as junk.
7. Select one of the four methods for dealing with junk mail messages from the "Junk mail messages should be" pop-up menu, depending on your organizations requirements.
   - Bounced: Sends the message back to the sender. You can send an email notification of the bounce to some email account, probably the postmaster.
   - Deleted: Deletes the message without delivery. You can send an email notification of the bounce to some email account, probably the postmaster.
   - Delivered: Delivers the message in spite of probably being junk mail. You can add text to the subject line, indicating that the message is probably junk mail or encapsulate the junk mail as a MIME attachment.
   - Redirected: Delivers the message to someone other than the intended recipient.
8. Optionally, choose to notify the intended recipient if the message was filtered in some way by attaching a subject tag to the email.
9. Choose how often to update the junk mail database. Enter the number in the "Update the junk mail and virus database __ time(s) every day." field.

   A minimum of twice a day is suggested. Some administrators choose eight times a day.
10. Click Save.

Enabling Virus Filtering

Mac OS X Server uses ClamAV to scan mail messages for viruses. If a suspected virus is found, you can choose to deal with it in several ways, as described below. The virus definitions are kept up to date (if enabled) via the Internet using a process called "freshclam." For more information about ClamAV, see www.clamav.net.

To enable virus filtering:

1. Open Server Admin.
2. In the Computers & Services list, select Mail for the server you are configuring.
3. Click Settings.
4. Select Filters.
5. Select "Scan email for viruses."
6. Select one of the three methods for dealing with infected mail messages from the "Infected messages should be" pop-up menu, depending on your organizations requirements.
   - Bounced: Sends the message back to the sender. You can send an email notification of the bounce to an email account (probably the domain's postmaster) and notify the intended recipient.
   - Deleted: Deletes the message without delivery. You can send an email notification to some email account, probably the postmaster, as well as the intended recipient.
   - Quarantined: Delivers the message to a folder for further analysis. You can send an email notification of the quarantine to some email account, probably the postmaster.
7. Optionally, choose to notify the intended recipient if the message was filtered in some way.
8. Choose how often to update the virus database. Enter the number in the "Update the junk mail and virus database __ time(s) every day." field.

   A minimum of twice a day is suggested. Some administrators choose eight times a day.
9. Click Save.

Disabling the SMTP Banner

The SMTP banner provides information about the mail server software running on the system, which could be useful to an attacker.

To replace the SMTP banner with a warning banner:

1. Open /etc/postfix/main.cf in a text editor.
2. Make sure any lines beginning with `smtpd_banner` are commented out and add the following line:

```
smtpd_banner = "Unauthorized use is prohibited."
```

Securing File Services

14

Securely configuring file services is an important step in the process of securing your private data against network attacks.

Mac OS X Server's cross-platform file sharing services help groups work more efficiently by letting them share resources, archive projects, exchange and back up important documents, and conduct other file-related activities.

Sharing files over a network opens your computers up to a host of vulnerabilities. With file services enabled, you are allowing access to files and folders on your server (also called share points).

For more information about configuring file services, see the file services administration guide.

Disabling File Services

Unless you are using the server as a file server, disable all file sharing services. Disabling these services prevents your computer from being used by an attacker to access other computers on your network.

To disable all file sharing services:

1. Open Server Admin.
2. Select the server you are configuring from the Computer & Services list.
3. Click AFP. Click Overview and verify that the pane says "Apple File Service is: Stopped." If not, click Stop.
4. Click FTP. Click Overview and verify that the pane says "FTP Service is: Stopped." If not, click Stop.
5. Click NFS. Click Overview and verify that the pane says "NFS Service is: Stopped."
6. Click Windows. Click Overview and verify that the pane says "Windows Service is: Stopped." If not, click Stop.

Choosing a File Sharing Protocol

If you require file sharing services, you must choose which file sharing protocols are needed before configuring your services. The protocol is configured for the folders you are sharing, called share points. The share points are created and configured using Workgroup Manager.

Most installations only need one file sharing protocol, and as few protocols as possible should be used. Limiting the number of protocols used by a server limits its exposure to vulnerabilities discovered in those protocols. The protocol choices are:

- Apple Filing Protocol (AFP)—AFP is the preferred method of file sharing for Macintosh or compatible client systems. AFP supports authentication of clients, and also supports encrypted network transport using SSH.
- File Transfer Protocol (FTP)—FTP should generally not be used for file sharing. The SFTP feature of the SSH protocol should be used instead. SFTP is designed to provide a secure means of authentication and data transfer, while FTP is not. The only situation where FTP is still an acceptable choice is when the server must act as a file server for anonymous users. This might be necessary over wide area networks, where there is no concern for the confidentiality of the data and responsibility for the integrity of the data rests with its recipient.
- Network File System (NFS)—NFS is a common file sharing protocol for UNIX computers. Avoid using NFS, because it does not perform authentication of its clients—it grants access based on client IP addresses and file permissions. Using NFS may be appropriate if the client computer administration and the network are trusted.
- Microsoft Windows Server Message Block (SMB/CIFS)—SMB is the native file sharing protocol for Microsoft Windows. Avoid using SMB—it supports authentication, but does not support encrypted network transport and it uses NTLMv1 and NTLMv2 encryption, both of which are very weak password hashing schemes. SMB may be an appropriate protocol for Windows clients when the network between the server and client is not at risk for eavesdropping.

Each of these protocols is appropriate for certain situations. Deciding which protocol to use depends on the clients and networking needs. After you choose a protocol for file sharing, you must configure the file sharing protocol.

Note: If no share points are shared with a particular protocol, then the service that runs that protocol can be disabled using the Server Admin program. The NFS service automatically stops when no share points specify its use.

Configuring AFP File Sharing Service

Apple File Service, which uses the Apple Filing Protocol (AFP), lets you share files among Macintosh clients. Because it provides both authentication and encryption, AFP service is the preferred file sharing method for Macintosh or compatible clients.

Note: Encryption does not apply to automatically mounted home folders, where only authentication is provided.

To securely configure AFP Service:

1. Open Server Admin.
2. In the Computers & Services list, select AFP for the server you are configuring.
3. Click Settings. Click General.
4. Deselect "Enable Bonjour registration."
5. Deselect "Enable browsing with AppleTalk."
6. Enter the login greeting according to site policy. Click Access.
7. For Authentication, choose "Kerberos" if your system is integrated into a Kerberos system. Otherwise, choose "Standard."
8. Deselect "Enable Guest access."
9. Select "Enable secure connections."
10. Deselect "Enable administrator to masquerade as any registered user."
11. Under Maximum Connections, enter the largest expected number for Client Connections.
12. Although Guest access was disabled, enter "1" for Guest Connections to minimize exposure in case it is accidentally reenabled. Click Logging.
13. Select "Enable access Log" to enable logging.
14. Select "Archive every __ day(s)." Set the frequency to three days or according to your organization's requirements.
15. Select "Login" and "Logout" to include those events in the access log. If you need stronger accounting, select the other events.
16. Under Error Log, select "Archive every __ day(s)." Set the frequency to three days or according to your organization's requirements.
17. Click Idle Users. The following Idle Users settings are suggested:
    - Deselect "Allow clients to sleep __ hour(s) - will not show as idle."
    - Select "Disconnect idle users after __ minute(s)" and enter a value in the text field to mitigate risk from a computer accidentally left unattended.
    - Deselect "Guests," "Administrators," "Registered users," and "Idle users" who have open files.
    - Enter a "Disconnect Message" notice according to site policy.

18 Click the "Start Service" button to begin using the file services.

For additional security enhancements, you can further restrict the AFP service by using SACLs and firewall rules. These are configured based on your organization's network environment.

- You can configure SACLs to restrict AFP access to specific users or groups. For more information, see "Setting Service Access Privileges" on page 188.
- You can configure firewall rules that prevent AFP connections from unintended sources. For more information, see "Creating Firewall Service Rules" on page 216.

Configuring FTP File Sharing Service

If authentication of users is possible, use the SFTP portion of the SSH protocol instead of the FTP server to securely transmit files to and from the server. For more information, see "Transferring Files Using SFTP" on page 197.

FTP is acceptable only if its anonymous access feature is required, which allows unauthenticated clients to download files. The files are transferred unencrypted over the network and no authentication is performed. Although the transfer does not guarantee confidentiality or integrity to the recipient, it may be appropriate in some cases. If this capability is not strictly required, disable it.

To configure the FTP to provide anonymous FTP downloads:

1 Open Server Admin.

2 In the Computers & Services list, select FTP for the server you are configuring.

3 Click Settings. Click General.

4 Enter 1 in "Disconnect client after __ login failures." Even though authenticated connections are not accepted, logins should fail quickly if accidentally activated.

5 Enter an email address specially set up to handle FTP administration—for example, ftpadmin@hostname.com.

6 Under Access, select "Kerberos" for Authentication. If a Kerberos server is not set up, the authentication process is blocked.

7 Enter 1 in "Allow a maximum of __ authenticated users." The GUI does not allow setting this to 0, but authenticated users are disabled in later steps.

8 Select "Enable anonymous access."

Anonymous access prevents the user credentials from being sent openly over the network.

9 Determine a maximum number of anonymous users and enter the number in "Allow a maximum of __ anonymous users."

10 Under File conversion, deselect "Enable MacBinary and disk image auto-conversion." Click Messages.

11 Select "Show Welcome Message" and enter a welcome message according to site policy.

12 Select "Show Banner Message" and enter a banner message in accordance with site policy. Do not reveal any software information, such as operating system type or version, in the banner. Click Logging.

13 Select all options under "Log authenticated users" and "Log anonymous users." Even though authenticated users are not allowed to log in, their attempts should be logged so corrective action can be taken. Click Advanced.

14 Set "Authenticated users see" to FTP Root and Share Points.

Both authenticated users and anonymous users see the same FTP root.

15 Verify that "FTP root" is set to the /Library/FTPServer/FTPRoot/ folder.

16 Click Save.

17 Open the /Library/FTPServer/FTPRoot/ folder and drag the contents (Users, Groups, Public) to the trash.

18 Drag the files you wish to share with anonymous users to the /Library/FTPServer/FTPRoot/ folder.

19 Verify that the file permissions for the /Library/FTPServer/FTPRoot/ folder do not allow public write access.

20 Open the file /Library/FTPServer/Configuration/ftpaccess for editing.

21 Delete any lines that begin with "upload." The following two line are present by default:

```
upload /Library/FTPServer/FTPRoot /uploads yes ftp daemon 0666 nodirs
upload /Library/FTPServer/FTPRoot /uploads/mkdirs yes ftp daemon 0666 dirs
     0777
```

22 Insert the following line to prevent advertisement of operating system and version information:

```
greeting terse
```

23 Insert the following lines to prevent any users from authenticating. This forces all users to access ftp anonymously, protecting their login credentials.

```
deny-gid %-99 %65535
deny-uid %-99 %65535
allow-gid ftp
allow-uid ftp
```

For additional security enhancements, you can further restrict the FTP service by using SACLs and firewall rules. These are configured based on your organizations network environment.

- You can configure SACLs to restrict FTP access to specific users or groups. For more information about configuring SACLs, see "Setting Service Access Privileges" on page 188.
- You can configure firewall rules that prevent FTP connections from unintended sources. For more information, see "Creating Firewall Service Rules" on page 216.

Configuring NFS File Sharing Service

NFS does not support user name and password authentication. It relies on client IP addresses to authenticate users, and on client enforcement of permissions that is not a secure approach in most networks. Therefore, use NFS only if you are on a local area network (LAN) with trusted client computers, or if you are in an environment that can't use Apple file sharing or Windows file sharing.

The NFS server included with Mac OS X Server lets you limit access to a share point based on a client's IP address. Access to a share point exported using NFS should be restricted to those clients that require it. You can reshare NFS mounts using AFP, Windows, and FTP so that users can access NFS volumes in a more restricted fashion.

To configure and start the NFS service, use Workgroup Manager. For information about how to setup and restrict NFS service, see "Configuring NFS Share Points" on page 133.

For additional security enhancements, you can further restrict the NFS service by using firewall rules. You can configure firewall rules that prevent AFP connections from unintended sources. For more information, see "Creating Firewall Service Rules" on page 216. These are configured based on your organization's network environment.

Configuring Windows File Sharing Service

If any share points are to use the SMB/CIFS protocol, then Windows file service must be activated and configured. Support for the SMB/CIFS protocol is provided by the open source Samba project, which is included with Mac OS X Server. SMB/CIFS uses NTLMv1 and NTLMv2 encryption, both of which are very weak password hashing schemes. For more information about configuring the Samba software, go to www.samba.org.

To securely configure Windows file service:

1. Open Server Admin.
2. In the Computers & Services list, select Windows for the server you are configuring.
3. Click Settings. Click General.
4. Choose the Role according to operational needs. If the server shares files, but does not provide authentication services, then "Standalone Server" is the appropriate choice.
5. Fill in the text fields appropriately. Leave the Description field blank. It is helpful for the computer name to match the host name (without the domain name). The Workgroup name depends on the configuration of Windows domains on your subnet.
6. Click Access. Deselect "Allow guest access."
7. For "Client connections" select "__ maximum," and enter the maximum number of client connections expected. The Graphs pane can display the actual usage, which can help you adjust the number appropriately for your network. Click Logging.
8. Change "Log Detail" to at least "medium" to capture authentication failures. Click Advanced.
9. Under Services, deselect "Workgroup Master Browser" and "Domain Master Browser" unless these services are operationally required.
10. Select Off for WINS registration.
11. Click Save.
12. Click the "Start Service" button to begin using the Windows service.

For additional security enhancements, you can further restrict the Windows service by using SACLs and firewall rules. These are configured based on your organizations network environment.

- You can configure SACLs to restrict Windows access to specific users or groups. For more information about configuring SACLs, see "Setting Service Access Privileges" on page 188.
- You can configure firewall rules that prevent Windows connections from unintended sources. For more information, see "Creating Firewall Service Rules" on page 216.

Restricting File Permissions

Before a folder is shared, its permissions should be restricted to the maximum extent possible. Permissions on share points set as user home folders are particularly important. By default, users' home folders are set to allow any other user to read their contents.

For more information about setting file permissions, see Chapter 6, "Securing Data and Using Encryption," on page 113.

Securing Web Service

15

Web service provides a great way to access data from anywhere in the world. However, this access is often attacked due to its weakness on other platforms. Mac OS X Server provides many rock-solid configuration options to protect and lock down web service.

Web technologies in Mac OS X Server are based on Apache, an open source HTTP web server. A web server responds to requests for HTML webpages stored on your site. Open source software allows anyone to view and modify the source code to make changes and improvements. This has led to Apache's widespread use, making it the most popular web server on the Internet today. Web administrators can use Server Admin to administer web technologies.

Fore more information about the Apache project, see www.apache.org. The Center for Internet Security (CIS) at www.cisecurity.org provides an Apache Benchmark and Scoring tool. CIS Benchmarks enumerate security configuration settings and actions that harden your computer.

For more information about configuring web service, see the web technologies administration guide.

Disabling the Web Service

Web server software should be deactivated if the system is not intended to be a web server. Secure web administration demands scrutiny of some basic configuration settings. SSL encryption should be used to encrypt any sensitive web traffic.

If the system is not intended to be a web server, disable web services using the Server Admin tool. Disabling the service prevents potential vulnerabilities on your computer. The web service is disabled by default, but verification is recommended.

To disable the web service:

1. Open Server Admin.
2. In the Computers & Services list, select Web for the server you are configuring.
3. Click Overview and verify that the pane says "Web Service is: Stopped." If not, click Stop Service.

Disabling Web Modules

If your system does not require active web modules they should be disabled. Web modules (sometimes called plug-ins) consist of web components that add functionality to the web service. Using unnecessary modules creates potential security risks when the web service is running.

There are many different web modules available for use with the web service. You should verify that each module used is required and that understand the impact it has to security when the web service is running.

Important: When disabling web modules make sure that the module is not needed by another web services your are running. If you disable a web module that another web services is dependent on that web service might not work.

To disable web modules:

1. Open Server Admin.
2. In the Computers & Services list, select Web for the server you are configuring.
3. Click Settings.
4. Click Modules.
5. Deselect all modules except for the modules that your site requires.
6. Click Save.

Disabling Web Options

The following web options should be deactivated unless they are specifically required for the web services. Activating these web option will automatically enable there associated web modules. This can be a security risk if you don't understand the impact the module has to security when the web service is running.

The following web modules should be deactivated unless they are specifically required for the web service:

- Folder Listing–Users can display a folder list of the website folder contents.
- WebDAV–WebDav lets you use a web server as a file server. Clients use their browsers from any location, on any type of computer, to access the shared files on the server.

- CGI Execution–Common Gateway Interface (CGI) scripts send information between your website and applications that provide different services for the site.
- WebMail–WebMail provides access to the mail service through a web browser.

To disable web options:

1. Open Server Admin.
2. In the Computers & Services list, select Web for the server you are configuring.
3. Click Settings.
4. Click Sites.
5. Select your site in the list and click the Edit (/) button. A new pane with configuration options for that site appears.
6. Click Options.

   Deselect "Folder Listing," "WebDAV," "CGI Execution," and "WebMail" unless they are required.

Configuring Web Service for SSL

If web service is required, protect the service communication with SSL. Enabling SSL for web service causes communications between the web server and its clients to be encrypted, protecting clients from eavesdroppers on the local network.

The steps to create SSL certificates for the web server are similar to those for the mail server. If the mail service and web service exist on the same server and use the same domain name, the same server certificate could be used for both services. However, this is not recommended.

The following steps describe the command-line method for creating certificates. For information about defining, obtaining, and installing certificates on your server using Certificate Manager in Server Admin, see "Readying Certificates" on page 174.

To securely configure web service with SSL:

1. Open Terminal and open the /usr/share/certs/ folder.

   If the /usr/share/certs folder does not exist create it.
2. Create a key pair for the web server using the `openssl` command.

   ```
   $ sudo openssl genrsa -des3 -out webserver.key 2048
   ```
3. When prompted, enter a strong, unique passphrase to protect the web server key pair.
4. Create the CSR with the web server key:

   ```
   $ sudo openssl req -new -key webserver.key -out webserver.csr
   ```

5 Enter the passphrase for the web server key pair and then fill out the following fields as completely as possible:

```
Country Name:
State or Province Name:
Locality Name (city):
Organization Name:
Organizational Unit Name:
Common Name:
Email Address:
```

The Common Name field is critically important. It must match the domain name of your server exactly (www.mypage.net) or the certificate will not work. Leave the challenge password and an optional company name blank.

6 Sign the webserver.csr certificate.

```
$ sudo openssl ca -in webserver.csr -out webserver.crt
```

7 When prompted, enter the CA passphrase to continue and then complete the process.

The certificate files needed to enable SSL on a web server are now located in the /usr/share/certs/ folder.

A separate certificate must be created for each domain name. For example, if a secure web page exists at www.mypage.net and a secure mail server is at mail.mypage.net, two certificates are needed. This is because the SSL protocol uses the certificate's Common Name field to verify the domain name.

8 Open Server Admin.

9 In the Computers & Services list, select Web for the server you are configuring.

10 Click Settings.

11 In the Sites pane, select your site in the list and click the Edit (/) button.

12 In the Security pane, select "Enable Secure Sockets Layer (SSL)."

When you turn on SSL, a message notes that the port is changed to 443.

13 Type the location of the SSL log file in the "SSL Log File" field.

You can also click the Browse (...) button to locate the folder you want to use. If you are administering a remote server, file service must be running on the remote server to use the Browse button.

14 Choose the certificate you want to use in the pop-up menu.

The name of the certificate must match the virtual host name if the certificate is protected by a passphrase. If the names don't match, web service won't restart.

If you choose Custom Configuration or want to edit a certificate, click the Edit (/) button and supply the correct information in each field for the certificate and click OK.

If you received a ca.crt file from the certificate authority, click the Edit button and paste the text from the ca.crt file in the Certificate Authority File field.

15 Click Save.

16 Confirm that you want to restart web service.

Using a Passphrase with SSL Certificates

Server Admin allows you to enable SSL with or without saving the SSL password. If you did not save the passphrase with the SSL certificate data, the server prompts you for the passphrase upon restart, but won't accept manually entered passphrases. Use the Security pane for the site in Server Admin to save the passphrase with the SSL certificate data.

If you manage SSL certificates using Server Admin and you use a passphrase for your certificates, Server Admin ensures that the passphrase is stored in the system keychain. When a website is configured to use the certificate and that web server is started, the `getsslpassphrase` utility extracts the passphrase from the system keychain and passes it to the web server, as long as the certificate name matches the virtual host name.

If you prefer not to rely on this mechanism, you can instead arrange for the Apache web server to prompt you for the passphrase when you start or restart it.

To prompt for a passphrase when starting and restarting Web service:

1 Open Terminal and enter the following command.

```
$ sudo serveradmin settings web:IfModule:_array_id:mod_ssl.c:SSL
     PassPhraseDialog=builtin
```

2 Start the Web service using the `serveradmin` tool.

```
$ sudo serveradmin start web
```

3 Enter the certificate passphrase when prompted.

Setting Up the SSL Log for a Website

SSL logging is important to security. After configuring SSL on your web server, you can set up a file to log SSL transactions and errors. This enhances security

To set up an SSL log:

1 Open Server Admin.

2 In the Computers & Services list, select Web for the server you are configuring.

3 Click Settings.

4 In the Sites pane, double-click the site you want to edit.

5 In the Security pane, select "Enable Secure Sockets Layer (SSL)."

6 Enter the pathname for the folder where you want to keep the SSL log in the "SSL Log File" field.

You can also use the Browse button to navigate to the folder.

7 Click Save.

Securing WebDAV

Web service includes support for Web-based Distributed Authoring and Versioning, known as WebDAV. With WebDAV capability, your client users can check out webpages, make changes, and then check the pages back in while the site is running. In addition, the WebDAV command set is rich enough that client computers with Mac OS X installed can use a WebDAV-enabled web server as if it were a file server.

Sharing files over a network opens your computers up to a host of vulnerabilities. To reduce the security risk when using WebDAV, assign access privileges for the sites and for the web folders.

To securely configure WebDAV for a site:

1. Open Server Admin.
2. In the Computers & Services list, select Web for the server you are configuring.
3. Click Settings.
4. In the Sites pane, select your site in the list and click the Edit (/) button.
5. In the Options pane, select WebDAV.
6. Click Save.
7. Click Realms. Select the realm and click Edit (/) or click the Add (+) button to create a new realm.

   The realm is the part of the website users can access.
8. Type the name you want users to see when they log in.
9. Choose the web authorization method from the Authorization pop-up menu. Basic authorization is on by default, but is not recommended.

   If you want digest authorization for the realm, choose "Digest." Digest authorization allows the user identity to be established without having to send a password in plaintext over the network.

   If you want Kerberos authorization for the realm, choose "Kerberos." To use Kerberos authorization for the realm, the server must be joined to a Kerberos domain and SSL must be on for the site. This is because credentials are sent in the clear, and Server Admin requires that SSL be on.
10. Type the path to the location in the website you want to limit access to, and click OK.

    You can also click the Browse (...) button to locate the folder you want to use.
11. Click Save when you have finished creating realms.

Note: If you have turned off the WebDAV module in the Modules pane of Server Admin, you must turn it on again before WebDAV takes effect for a site. This is true even if the WebDAV option is checked in the Options pane for the site.

Setting Access for Websites

Web content provided by the web server should be protected by setting user and group permissions. This restricts access to the web content (files and folders) to the authors. Web content is stored by default in the /Library/WebServer/Documents/ folder.

Use realms to control access and provide security for websites by specifying who has access to them. Realms are locations within a site (or the site itself) that users can view. If WebDAV is enabled, users who have authoring privileges can also make changes to content in the realm. You set up the realms and specify which users and groups have access to them.

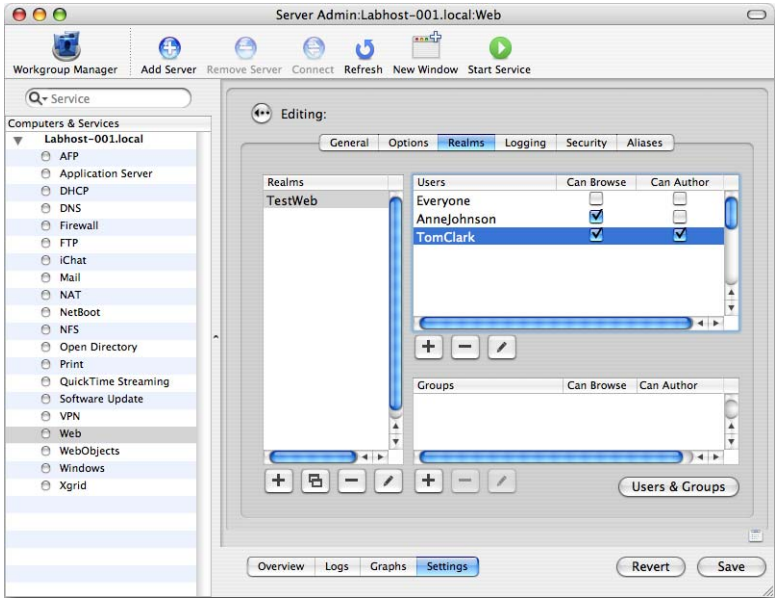

To set web access using a realm:

1 Open Server Admin.

2 In the Computers & Services list, select Web for the server you are configuring.

3 Click Settings.

4 In the Sites pane, select your site in the list and click the Edit (/) button.

5 In the Realms pane, select the realm you want to edit.

If no realm names are listed, create one.

6 To specify access for individual users and groups, click Users & Groups to open a drawer listing users and groups.

7 Click Users or Groups in the drawer's button bar to show the list you want.

8 Drag user names to the Users field or group names to the Groups field.

Note: You can also click the Add (+) button to open a sheet in which you type a user or group name and select access options.

9 Select Can Browse and/or Can Author for each user and group, as required.

Can Browse allows anyone with access to this realm browsing capability.

Can Author allows anyone with access to this realm to make changes to the web content.

Can Browse and Can Author together allows anyone with access to this realm to see and make changes to it.

Note: When users or members of a group you've added to the realm connect to the site, they must supply their user name and password.

10 Click Save.

Securing Weblogs

Blogging, the practice of using a webpage as an electronic journal or newsletter, has become a popular way to exchange information among users with common interests. The webpage, referred to as a weblog or a blog, is used to post entries and display them in chronological order. Because they're web-based, entries can include electronic links, which make viewing related content fast and easy.

By default, weblogs are disabled when you start the web service. Running weblogs can open your computers up to a host of vulnerabilities. If weblogs are not required, they should be disabled.

Enabling Weblogs

The web server provides weblogs (blogs) as an option for each website. If you require weblogs, you must enable web service to provide SSL communication between the users and the server.

Each server can host one weblog for each user and group defined in folders in the server's Open Directory search path. The weblogs comply with RSS and Atom XML standards and allow Open Directory authentication.

You activate weblogging by starting web service in Server Admin, then establishing a few default weblog settings. You can set up one Weblog server per server computer.

Note: When you turn on weblogs, they are on for every site on the server.

To set up Weblog service:

1. Open Server Admin.
2. In the Computers & Services list, select Web for the server you are configuring.
3. If web service isn't running, click Start Service.

   To maximize the security of user interactions with the server hosting weblogs, have users access weblogs through a site that has SSL enabled.
4. Click Settings.
5. Click Weblogs.
6. Select Enable Weblogs. This setting enables weblog access through any website that web service is configured to support.
7. Choose a default theme.

   A theme controls the appearance of a weblog. Themes determine the color, size, location, and other attributes of weblog elements. Each theme is implemented using a style sheet.

   The default theme is used when a weblog is initially created, but weblog owners can change the theme any time. The default theme also controls the appearance of Weblog service's front page.
8. Identify a weblog folder, used to store weblog files.

   By default, weblog files are stored in /Library/Application Support/Weblogs/ on the computer hosting Weblog service. You can click Browse to select a different folder, such as a folder on a RAID device or on another computer.
9. Optionally, specify a default email domain, such as "example.com."

   The email domain is used to construct an initial email address for a weblog. For example, a user whose short name is "anne" would have an email address of anne@example.com displayed on his weblog initially, but the user can change it.
10. Click Save.
11. Make sure that the Weblog server's Open Directory search path includes directories in which users and group members you want to support with Weblog service are defined.

    The Open Directory administration guide explains how to set up search paths. Any user or group member defined in the Open Directory search path is now authorized to create and access weblogs on the server unless you deny them access to Weblog service.

Securing the Application Server

JBoss (version 3.2.3) is an open source application server designed for J2EE applications. It runs on Java 1.4.2. JBoss is a widely used, full-featured Java application server. It provides a full Java 2 Platform, Enterprise Edition (J2EE) technology stack with features, such as:

- An Enterprise Java Bean (EJB) container
- Java Management Extensions (JMX)
- Java Connector Architecture (JCA)

By default, JBoss uses Tomcat as its web container, but you can use other web containers, such as Jetty, if you wish. Additional information about these Java technologies is available online.

- For JBoss, see www.jboss.org/.
- For J2EE, see java.sun.com/j2ee/.

For more information about configuring the application server and how to deploy and manage J2EE applications using JBoss, see the Java application server administration guide.

Disabling the Application Server

The application server software should be disabled if your server is not intended to be an application server. Disabling the service prevents potential vulnerabilities on your computer. The web service is disabled by default, but verification is recommended.

To disable the application server:

1. Open Server Admin.
2. In the Computers & Services list, select Application Server for the server you are configuring.
3. Click Settings.
4. Click Overview and verify that the pane says "Application Server is: Stopped." If not, click Stop Service.

Securely Configuring the Application Server

If application server software is required, use Server Admin to enable it. With JBoss turned on, you can use the management tool to configure your server.

To start the application server:

1. Open Server Admin
2. In the Computers & Services list, select Application Server for the server you are configuring.
3. Click Settings.
4. Click General.
5. Select one of the JBoss options.

   JBoss is preconfigured to use a local configuration.

   With JBoss turned on, you can use the management tool to configure your server.

Backing Up and Restoring Application Server Configurations

Use Server Admin to back up and restore JBoss configurations.

To back up your application server configuration:

1. Open Server Admin.
2. In the Computers & Services list, select Application Server for the server you are configuring.
3. Click Settings.
4. Click Backup.
5. Click Backup and navigate to the location where you want to store configurations.

Securing WebObjects

Mac OS X Server includes the WebObjects runtime libraries and an unlimited deployment license, making it the ideal platform for your J2EE-compatible WebObjects applications.

You can optionally purchase the WebObjects development tools from the Apple Store (store.apple.com), Apple's retail stores, and authorized Apple resellers.

Disabling WebObjects

The WebObjects software should be disabled if your server is not intended to be an WebObjects server. Disabling the service prevents potential vulnerabilities on your computer. The WebObjects service is disabled by default, but verification is recommended.

To disable the WebObjects service:

1. Open Server Admin.
2. In the Computers & Services list, select WebObjects for the server you are configuring.
3. Click Overview and verify that the pane says "WebObject Service is: Stopped." If not, click Stop Service.

Securing Client Configuration Management Services

16

Securely configuring client configuration management helps standardize the clients across your network and provide a secure deployment.

Network computers can be managed through NetBoot, which decreases maintenance time and can help to protected against malicious software attacks. You can further protect against attacks by configuring an internal Software Update server. This allows you to maintain a secure network by controlling what software updates are installed on your network computers.

Securing NetBoot Service

Using NetBoot you can have your client computers start up from a standardized Mac OS X configuration suited to their specific tasks. Because the client computers start up from the same image, you can quickly update the operating system for an entire group by updating a single boot image.

A boot image is a file that looks and acts like a mountable disk or volume. NetBoot boot images contain the system software needed to act as a startup disk for client computers via the network. An install image is a special boot image that starts the client long enough to install software from the image, after which the client can start up from its own hard disk. Both boot images and install images are special kinds of disk images. Disk images are files that behave just like disk volumes.

For more information about configuring NetBoot service, see the system image administration guide.

Disabling NetBoot Service

The NetBoot service should be disabled if your server is not intended to be a NetBoot server. Disabling the service prevents potential vulnerabilities on your computer. The NetBoot service is disabled by default, but verification is recommended.

The best way to prevent clients from using NetBoot on the server is to disable NetBoot service on all Ethernet ports.

To disable NetBoot:

1. Open Server Admin.
2. In the Computers & Services list, select NetBoot for the server you are configuring.
3. Click Overview and verify that the pane says "NetBoot Service is: Stopped." If not, click Settings.
4. Click General.
5. Deselect Enable for any ports listed.
6. Click Save.

Securely Configuring NetBoot Service

If NetBoot service is required, securely configure it with restrictions on the ports it uses, the images available, and client access to the service. NetBoot service uses AFP, NFS, DHCP, Web, and TFTP services, depending on the types of clients your are trying to boot. These other services must also be configured securely to reduce network vulnerabilities.

NetBoot creates share points for storing boot and install images in /Library/NetBoot/ on each volume you enable and names them NetBootSP*n*, where *n* is 0 for the first share point and increases by 1 for each additional share point. If, for example, you decide to store images on three separate server disks, NetBoot sets up three share points named NetBootSP0, NetBootSP1, and NetBootSP2.

The share points for client shadow files are also created in /Library/NetBoot/ and are named NetBootClients*n*.

The initial NetBoot process relies on simple hardware authentication by validating the hardware address of the client's computer's network interface accessing the server prior to receiving the preconfigured system image. For that reason, maintain and use these services over a trusted network.

To securely configure NetBoot:

1. Open Server Admin.
2. In the Computers & Services list, select NetBoot for the server you are configuring.
3. Click Settings.
4. Click General, and select which network ports to use for providing NetBoot service.

   This enables one or more network ports to serve NetBoot images.
5. Click Images, and select the boot or install images to serve.

   Only enable the images that are approved for use in your environment. The best way to prevent clients from using NetBoot disk images is to disable them. Disabling an image prevents client computers from starting up using the image.

6 Click Filters.

NetBoot service filtering lets you restrict access to the service based on the client's Ethernet hardware address (MAC). A client's address is added to the filter list automatically the first time it starts up from an image on the server and is allowed access by default.

7 Select Enable NetBoot filtering.

8 Select either "Allow only clients listed below" or "Deny only clients listed below."

9 Use the Add (+) and Delete (–) buttons to modify the list of clients.

To look up a MAC address, type the client's DNS name or IP address in the Host Name field and click Search.

To find the hardware address for a computer using Mac OS X, look on the TCP/IP pane of the computer's Network preference or run Apple System Profiler.

10 Click Save.

Viewing NetBoot Service Logs

NetBoot service logging is important to security. With logs, you can monitor and track client communication to the NetBoot server. The NetBoot service log is /var/log/system.log that can be accessed using Server Admin.

To view the NetBoot service log:

1 Open Server Admin.

2 In the Computers & Services list, select NetBoot for the server you are logging

3 Click Settings.

4 Click Logging.

You also have the option of filtering the log details by choosing Low, Medium, or High from the Log Detail Level pop-up menu.

5 Click Save.

Securing Software Update Service

Software Update service lets you manage Macintosh software updates from Apple on your network. In an uncontrolled environment, users can connect to the Apple Software Update servers at any time and update their computer with software that is not approved for use by your organization. By configuring a Software Update server you can make sure that only approved software updates are installed on your network computers.

Disabling Software Update Service

The Software Update service should be disabled if your server is not intended to be a software update server. Disabling the service, prevents potential vulnerabilities on your computer. The software update service is disabled by default, but verification is recommended.

To disable software update service:

1. Open Server Admin.
2. In the Computers & Services list, select Software Update for the server you are configuring.
3. Click Overview and verify that the pane says "Software Update Service is: Stopped." If not, click Settings.
4. Click General.
5. Deselect Enable for any ports listed.
6. Click Save.

Securing Directory Services

17

Directory services are the backbone of your network's security policy. The granting of access to the information and services on your network should be well planned and thought out.

A directory service provides a central repository for information about computer users and network resources in an organization. Mac OS X Server's Open Directory provides directory and authentication services for mixed networks of Mac OS X, Windows, and UNIX computers. Open Directory uses OpenLDAP, the open source implementation of the Lightweight Directory Access Protocol (LDAP), to provide directory services.

Open Directory also provides authentication service. It can securely store and validate the passwords of users who want to log in to client computers on your network or use other network resources that require authentication. Open Directory can also enforce policies such as password expiration and minimum length.

For more information about passwords and authentication, see Appendix A, "Understanding Passwords and Authentication," on page 295.

The Open Directory service must be set to the proper role and configured to use SSL to encrypt its communications to protect the confidentiality of its important authentication data. Password policies can also be enforced by the Open Directory service.

For more information about understanding and configuring directory and authentication services, see the Open Directory administration guide.

Understanding Open Directory Server Roles

The Open Directory services can be configured to one of several roles, depending on the server's place in the overall network and directory structure:

- Standalone Server—This role does not share information with other computers on the network. It is a local directory domain only.
- Connected to a Directory System—This role allows the server to get directory and authentication information from another server's shared directory domain.
- Open Directory Master—This role provides an Open Directory Password Server, which supports all conventional authentication methods required by Mac OS X Server services. In addition, an Open Directory Master can provide Kerberos authentication for single sign-on.
- Open Directory Replica—This role acts as a backup to the Open Directory master. It can provide the same directory and authentication information to other networks as the master. It has a read-only copy of the master's LDAP directory domain.

Configuring the Open Directory Services Role

If the server is not intended to be a directory server, make sure the LDAP server is stopped using Server Admin. To stop LDAP server, set the Open Directory role to Standalone Server. This prevents Open Directory from engaging in unnecessary network communications. On a newly installed server, the LDAP server should be stopped by default, but verification is recommended.

To configure the Open Directory role:

1. Open Server Admin.
2. In the Computers & Services list, select Open Directory for the server you are configuring.
3. Click Settings.
4. Click General.
5. Choose the role for your server from the Role pop-up menu.

   Set the role to Standalone Server if Open Directory services are not required.

   If an Open Directory Master role is required, make sure that only legitimate replicas are listed and replicate to clients whenever the directory is modified.

   If an Open Directory Replica role is required, make sure that the intended master is set.

   If a Connected to a Directory System role is required, make sure that the Open Directory has joined the appropriate Kerberos realm.
6. Click Save.

Starting Kerberos After Setting Up an Open Directory Master

Kerberos was designed to solve network security problems. It does not transmit the user's password across the network, or save it in the user's computer memory or disk. This keeps an attacker from learning the original password if they compromise or crack the Kerberos credentials, allowing an attacker to potentially compromise only a small portion of the network rather than the whole network.

If Kerberos doesn't start automatically when you set up an Open Directory master, you can use Server Admin to start it manually. First you have to fix the problem that prevented Kerberos from starting. Usually the problem is DNS service that isn't configured correctly or isn't running at all.

Note: After you manually start Kerberos, users whose accounts have Open Directory passwords and were created in the Open Directory master's LDAP directory while Kerberos was stopped may have to reset their passwords the next time they log in. A user account is thus affected only if all the recoverable authentication methods for Open Directory passwords were disabled while Kerberos was stopped.

To start Kerberos manually on an Open Directory master:

1. Open Server Admin.
2. In the Computers & Services list, select Open Directory for the server you are configuring.
3. Click Refresh (or choose View > Refresh) and check the status of Kerberos as reported in the Overview pane.

   If Kerberos is running, there's nothing more to do.
4. Use Network Utility (in /Applications/Utilities/) to do a DNS lookup of the Open Directory master's DNS name and a reverse lookup of the IP address.

   If the server's DNS name or IP address don't resolve correctly:

   - In the Network pane of System Preferences, look at the TCP/IP settings for the server's primary network interface (usually built-in Ethernet). Make sure the first DNS server listed is the one that resolves the Open Directory server's name.
   - Check the configuration of DNS service and make sure it's running.
5. In Server Admin, select Open Directory for the server you are configuring.
6. Click Settings.

7 Click General.

8 Click Kerberize, then enter the information requested.

Administrator Name and *Password*: You must authenticate as an administrator of the Open Directory master's LDAP directory.

Realm Name: This field is preset to be the same as the server's DNS name converted to capital letters. This is the convention for naming a Kerberos realm. You can enter a different name if necessary.

Configuring Open Directory for SSL

Using Server Admin, you can enable SSL for encrypted communications between an Open Directory server's LDAP directory domain and computers that access it. SSL uses a digital certificate to provide a certified identity for the server. You can use a self-signed certificate or a certificate obtained from a certificate authority. Generating SSL certificates for LDAP services is similar to generating SSL certificates for the web server.

SSL communications for LDAP use port 636. If SSL is disabled for LDAP service, communications are sent as clear text on port 389.

The following steps describe the command-line method for creating certificates. For information about defining, obtaining, and installing certificates on your server using Certificate Manager in Server Admin, see "Readying Certificates" on page 174.

To create a new Open Directory service certificate:

1 Generate a private key for the server in the /usr/share/certs/ folder:

If the /usr/share/certs folder does not exist create it.

```
$ sudo openssl genrsa -out ldapserver.key 2048
```

2 Generate a CSR for the CA to sign:

```
$ sudo openssl req -new -key ldapserver.key -out ldapserver.csr
```

3 Fill out the following fields as completely as possible, making certain that the Common Name field matches the domain name of the LDAP server exactly:

```
Country Name:
Organizational Unit:
State or Province Name:
Common Name:
Locality Name (city):
Email Address:
Organization Name:
```

Leave the challenge password and optional company name blank.

4 Sign the ldapserver.csr request with the `openssl` command.

```
$ sudo openssl ca -in ldapserver.csr -out ldapserver.crt
```

5 When prompted, enter the CA passphrase to continue and complete the process.

The certificate files needed to enable SSL on the LDAP server are now in the /usr/share/certs/ folder.

6 Open Server Admin.

7 In the Computers & Services list, select Open Directory for the server that is an Open Directory master or an Open Directory replica.

8 Click Settings.

9 Click Protocols.

10 Choose "LDAP Settings" from the Configure pop-up menu.

11 Select Enable Secure Sockets Layer (SSL).

12 Use the Certificate pop-up menu to choose an SSL certificate that you want LDAP service to use.

The menu lists all SSL certificates that have been installed on the server. To use a certificate not listed, choose Custom Configuration from the pop-up menu.

13 Click Save.

Configuring Open Directory Policies

You can set password, binding, and security policies for an Open Directory master and its replicas. You can also can set several LDAP options for an Open Directory master or replica.

For more information about configuring policies, see "Configuring User Accounts" on page 136.

Setting the Global Password Policy

Using Server Admin, you can set a global password policy for user accounts in a Mac OS X Server directory domain. The global password policy affects user accounts in the server's local directory domain. If the server is an Open Directory master or replica, the global password policy also affects user accounts that have an Open Directory password type in the server's LDAP directory domain. If you change the global password policy on an Open Directory replica, the policy settings are eventually synchronized with the master and any other replicas of it.

Administrator accounts are always exempt from password policies. Each user can have an individual password policy that overrides some of the global password policy settings.

Kerberos and Open Directory Password Server maintain password policies separately. Mac OS X Server synchronizes the Kerberos password policy rules with Open Directory Password Server password policy rules. Replicas of the Open Directory master automatically inherit its global password policy.

To change the global password policy of all user accounts in the same domain:

1. Open Server Admin.
2. In the Computers & Services list, select Open Directory for the server you are configuring.
3. Click Settings.
4. Click Policy.
5. Click Passwords.

   This allows you to set the password policy options you want enforced for users who do not have their own individual password policies.
6. Select "differ from account name."
7. Select "contain at least one letter."
8. Select "contain at least one numeric character."
9. Select "be reset on first user login."
10. Select "contain at least 12 characters."
11. Select "differ from last 3 passwords used."
12. Select "be reset every 3 months."

    Note: If you select an option that requires resetting the password, remember that some service protocols don't allow users to change passwords. For example, users can't change their passwords when authenticating for IMAP mail service.
13. Click Save.

Setting a Binding Policy for an Open Directory Master and Replicas

Using Server Admin, you can configure an Open Directory master to allow or require trusted binding between the LDAP directory and the computers that access it. Replicas of the Open Directory master automatically inherit its binding policy.

Trusted LDAP binding is mutually authenticated. The computer proves its identity by using an LDAP directory administrator's name and password to authenticate to the LDAP directory. The LDAP directory proves its authenticity by means of an authenticated computer record that's created in the directory when you set up trusted binding.

Clients can't be configured to use both trusted LDAP binding and a DHCP-supplied LDAP server (also known as DHCP option 95). Trusted LDAP binding is inherently a static binding, but DHCP-supplied LDAP is a dynamic binding.

Note: Clients need version 10.4 or later of Mac OS X or Mac OS X Server to use trusted LDAP binding. Clients using version 10.3 or earlier are not able to set up trusted binding.

To set the binding policy for an Open Directory master:

1. Open Server Admin.
2. In the Computers & Services list, select Open Directory for the server that is an Open Directory master.
3. Click Settings (near the bottom of the window).
4. Click Policy (near the top).
5. Click Binding, and then set the directory binding options you want.

   To allow trusted binding, select "Enable directory binding."

   To require trusted binding, select "Require clients to bind to directory."
6. Click Save.

Setting a Security Policy for an Open Directory Master and Replicas

Using Server Admin, you can configure a security policy for access to the LDAP directory of an Open Directory master. Replicas of the Open Directory master automatically inherit its security policy.

To set the security policy for an Open Directory master:

1. Open Server Admin.
2. In the Computers & Services list, select Open Directory for the server that is an Open Directory master.
3. Click Settings.
4. Click Policy.
5. Click Binding, ard then set the security options you want.

   "Disable clear text passwords" determines whether clients can send passwords as clear text if the passwords can't be validated using any authentication method that sends an encrypted password.

   "Digitally sign all packets (requires Kerberos)" ensures directory data from the LDAP server won't be intercepted and modified by another computer while en route to client computers.

   "Encrypt all packets (requires SSL or Kerberos)" requires the LDAP server to encrypt directory data using SSL or Kerberos before sending it to client computers.

   "Block man-in-the-middle attacks (requires Kerberos)" protects against a rogue server posing as the LDAP server. This option is best if used with the "Digitally sign all packets" option.
6. Click Save.

Securing Print Service

18

Print service is often an over-looked part of a security configuration. Important information passes into your networked printers and it is important that your printers are not misused.

With a print server, you can share printers by setting up print queues accessible by any number of users over a network connection. When a user prints to a shared queue, the print job waits on the server until the printer is available or until established scheduling criteria are met.

Apple's printing infrastructure is built on Common UNIX Printing System (CUPS). CUPS uses open standards such as Internet Printing Protocol (IPP) and PostScript Printer Description files (PPDs).

For more information about configuring print service, see the print service administration guide.

Disabling Print Service

The print server software should be disabled if your server is not intended to be a print server. Disabling the service prevents potential vulnerabilities on your computer. The print service is disabled by default, but verification is recommended.

To disable print service:

1. Open Server Admin.
2. In the Computers & Services list, select Print for the server you are configuring.
3. Click Overview and verify that the pane says "Print Service is: Stopped." If not, click Stop Service.

Configuring Print Queues

If print service is required, you should create a print queue for shared printers that is accessible by users over a network connection.

AppleTalk and Line Printer Remote (LPR) printer queues do not support authentication. Print service relies on the client to provide user information. Although standard Macintosh and Windows clients provide correct information, a clever user could potentially modify the client to submit false information and thereby avoid print quotas.

Windows service does support authentication, requiring users to log in before using SMB/CIFS printers. Print service uses both Basic and Digest (MD5) authentication and supports a print job submission method called Internet Printing Protocol (IPP).

To create a print queue:

1. Open Server Admin.
2. In the Computers & Services list, select Print for the server you are configuring.
3. Click Settings.
4. Click Queues.
5. Click the Add (+) button to create a new queue.
6. Choose the protocol used by the printer from the pop-up menu.

   For an AppleTalk printer, select the printer in the list and click OK.

   For an LPR printer, type the printer IP address or DNS name and click OK. If you don't want to use the printer's default queue, first deselect "Use default queue on server" and type a queue name.

   For an Open Directory printer, select the printer in the list and click OK.
7. In the Queues pane, select the queue you have just added and click the Edit (/) button.
8. Type the queue name you want clients to see in the Sharing Name field.

   Make sure the name is compatible with any naming restrictions imposed by your clients. For example, some LPR clients do not support names that contain spaces, and some Windows clients restrict names to 12 characters.

   Queue names shared via LPR or SMB/CIFS should not contain characters other than A–Z, a–z, 0–9, and _ (underscore).

   AppleTalk queue names cannot be longer than 32 bytes (which may be fewer than 32 typed characters). Note that the queue name is encoded according to the language used on the server and may not be readable on client computers using another language.

   Note: Changing the Sharing Name also changes the queue name that appears in Printer Setup Utility on the server.

9 Select the protocols used for printing by your client computers. If you select "SMB", make sure you start Windows services.

10 Select "Enforce quotas for this queue" if you want to enforce the print quotas you establish for users in Workgroup Manager.

11 Select a cover sheet you would like the printer to create by selecting a title from the pop-up menu. If you don't want the server to create a cover sheet for each job printed, select "none".

12 Click Save.

Configuring Print Banners

Print banners protect confidential information while it is printing. Print banners are used to separate print jobs as they print on a network printer. They also can include information that identifies the printed materials as confidential information.

Creating Banner Pages

To create a print banner page, copy the file you are using as the banner to the /usr/share/cups/banners/ folder. Banner files can be any file format supported by CUPS and can be any number of pages. Typically banner pages are ASCII text and fit on one side of a page. The standard CUPS banner pages are single-page PostScript files.

CUPS includes the following banner files:

- None—Do not produce a banner page.
- Classified—A banner page with a "classified" label at the top and bottom.
- Confidential—A banner page with a "confidential" label at the top and bottom.
- Secret—A banner page with a "secret" label at the top and bottom.
- Standard—A banner page with no label at the top and bottom.
- Topsecret—A banner page with a "top secret" label at the top and bottom.
- Unclassified—A banner page with an "unclassified" label at the top and bottom.

To save a printer banner option to a printer, use the lpoptions command:

```
$ lpoptions -p laserjet -o job-sheets=classified filename
```

Note: Running the `lpoptions` command as the root user sets the default options for all users. The root account does not have its own set of default options.

Securing Multimedia Services

19

Protecting QuickTime multimedia streams and only allowing access to those who are authorized to view can help keep information private. The following section will help you better understand and configure QTSS securely.

Streaming is the delivery of media, such as movies and live presentations, over a network in real time. A computer (streaming server) sends the media to another computer (client computer), which plays the media as it is delivered.

With QuickTime Streaming Server (QTSS) software, you can deliver:

- Broadcasts of live events in real time
- Video on demand
- Playlists of prerecorded content

A certain level of security is inherent in real-time streaming, since content is delivered only as the client needs it and no files remain afterward, but you might need to address some security issues.

For more information about configuring multimedia services, see the QuickTime Streaming Server administration guide.

Disabling QuickTime Streaming Server

The QuickTime Streaming server software should be disabled if your server is not intended to be a QuickTime streaming server. Disabling the service, prevents potential vulnerabilities on your computer. The QuickTime Streaming service is disabled by default, but verification is recommended.

To disable streaming service:

1. Open Server Admin.
2. In the Computers & Services list, click QuickTime Streaming for the server you are configuring.
3. Click Overview and verify that the pane says "QuickTime Streaming Server is: Stopped." If not, click Stop Service.

Configuring a Streaming Server

If you require QuickTime streaming service, configure it in conjunction with your firewall and bind it to a single IP address.

To configure a streaming server:

1. Open Server Admin.
2. In the Computers & Services list, click QuickTime Streaming for the server.
3. Click Settings.
4. Click IP Binding.

   By binding QuickTime Streaming Server with an IP address, you can easily track network activity. You can also configure the Firewall to restrict network access to this IP address. IP binding is also helpful when your server is multihomed (for example, if you're also hosting a web server).
5. Select the IP address from the list.
6. Click Save.
7. Start Service.

Controlling Access to Streamed Media

You can set up authentication to control client access to streamed media files. Two schemes of authentication are supported: basic and digest. By default, the server uses the more secure digest authentication.

You can also control playlist access and administrator access to your streaming server. Authentication does not control access to media streamed from a relay server. The administrator of the relay server must set up authentication for relayed media.

The ability to manage user access is built into the streaming server, so it is always enabled. For access control to work, an access file must be present in the folder you selected as your media directory. If an access file is not present in the streaming server media directory, all clients are allowed access to the media in the folder.

Users must have QuickTime 5 or later to access a media file for which digest authentication is enabled. If your streaming server is set up to use basic authentication, users need QuickTime 4.1 or later. Users must enter their user name and password to view the media file. Users who try to access a media file with an earlier version of QuickTime will see the error message "401: Unauthorized."

To set up access control:

1. Open Terminal.
2. Add new user accounts and passwords to the QuickTime Streaming service.

```
$ sudo qtpasswd /Library/QuickTimeStreaming/Config/qtusers annejohnson
```

3. Enter a password for the user and reenter it when prompted.
4. Create an access file and place it in the media directory that you want to protect.
5. If you want to disable authentication for a media directory, remove the access file (called qtaccess) or rename it (for example, qtaccess.disabled).

Creating an Access File

An access file is a text file called "qtaccess" that contains information about users and groups who are authorized to view media in the folder in which the access file is stored. The folder you use to store streamed media can contain other folders, and each folder can have its own access file. When a user tries to view a media file, the server checks for an access file to see whether the user is authorized to view the media. The server looks first in the folder where the media file is located. If an access file is not found, it looks in the enclosing folder. The first access file that's found is used to determine whether the user is authorized to view the media file.

To create an access file:

1. Open Terminal.
2. Create an access file with a text editor such as `vi` or `pico`.

```
$ vi qtaccess
```

3. Add the following information to the file:

```
AuthName message
AuthUserFile user filename
AuthGroupFile group filename
require user username1 username2
require group groupname1 groupname2
require valid-user
require any-user
```

The following table describes each of the variables you can use when customizing your access file.

Variable	Description
message	Text your users see when the login window appears. It's optional. If your message contains any white space (such as a space character between terms), make sure you enclose the entire message in quotation marks.
user filename	The path and filename of the user file. The default is /Library/QuickTimeStreaming/Config/qtusers.
group filename	The path and filename of the group file. The default is /Library/QuickTimeStreaming/Config/qtgroups. A group file is optional. If you have a lot of users, it may be easier to set up one or more groups, then enter the group names, than to list each user.
username	A user who is authorized to log in and view the media file. The user's name must be in the user file you specified. You can also specify "valid-user," which designates any valid user.
groupname	A group whose members are authorized to log in and view the media file. The group and its members must be listed in the group file you specified.
valid-user	Any user defined in the qtusers file. The statement "require valid-user" specifies that any authenticated user in the qtusers file can have access to the media files. If this tag is used, the server will prompt users for an appropriate user name and password.
any-user	Allows any user to view media without providing a name or password.
AuthScheme	Set with the values "basic" or "digest" to override the global authentication setting on a directory-by-directory basis.

4 Save the access file and exit the editor.

```
$ :wq
```

Setting Up Relay Streams

You use relays to accept a stream from one streaming server and send the stream on, or "relay" it, to another streaming server. Each relay comprises a source and one or more destinations.

To set up a relay:

1 In the Settings pane of the QuickTime Streaming service, click Relays.

2 Click the Add (+) button next to the Relays list.

3 Enter a name for the relay in the Relay Name field.

4 Choose an option from the Relay Type pop-up menu.

The Relay Type defines the source for the relay. There are three options for Relay Type:

- **Request Incoming Stream** directs the streaming server to send a request to the source computer for the incoming stream before it gets relayed. You can use this feature to relay a reflected live broadcast (from another server) or to request a stored file and turn it into an outgoing live stream. Request Incoming Stream is commonly used with unannounced UDP streams from QuickTime Broadcaster or other streaming encoders.
- **Unannounced UDP** directs the server to relay streams on a specific IP address and port numbers.
- **Announced UDP** directs the server to wait for the incoming stream and then relay it. Relays set to wait for announced streams can accept only media streams using the RTSP announce protocol. Announced UDP is used with Automatic (Announced) broadcasts from QuickTime Broadcaster or other streaming encoders that support the RTSP announce protocol.

5 In the case of Request Incoming Stream or Announced UDP, in the Source IP field, enter the DNS host name or IP address of the source computer.

6 In the case of Request Incoming Stream or Announced UDP, in the Path text field, enter the pathname to the stream.

7 In the case of Request Incoming Stream or Announced UDP, if the source computer requires automatic broadcasts to be authenticated, enter a user name and password.

8 Make sure Enable Relay is selected and click the Back button.

9 Click the Add (+) button next to the Destinations list.

There are two types of destinations:

- **Unannounced UDP** directs the server to relay the stream to the destination IP address and port numbers. This requires generating a session description protocol (sdp) file manually.
- **Announced UDP** directs the server to relay and announce the stream to the destination IP address. The sdp file will be generated automatically at the destination.

10 Enter the requested information and click the Back button.

11 Repeat steps 9 and 10 for each destination, and then click Save.

To turn a relay on or off, select or deselect the Enable checkbox next to the relay in the list. To delete a relay, select it and click the Delete (–) button.

Securing Grid and Cluster Computing Services

20

Protecting grid and clustering services helps control your network's free CPU cycles from misuse. This chapter helps you restrict your network's CPUs to authorized users.

Xgrid, a technology in Mac OS X Server and Mac OS X, simplifies deployment and management of computational grids. Xgrid enables you to group computers into grids or clusters, and allows users to easily submit complex computations to groups of computers (local, remote, or both), as either an ad hoc grid or a centrally managed cluster.

For more information about configuring multimedia services, see the Xgrid administration guide.

Disabling Xgrid Service

The Xgrid server software should be disabled if your server is not intended to be a Xgrid server. Disabling the service prevents potential vulnerabilities on your computer. The Xgrid service is disabled by default, but verification is recommended.

To disable the Xgrid service:

1. Open Server Admin.
2. In the Computers & Services list, select Xgrid for the server that you are configuring.
3. Click Overview and verify that the pane says "Xgrid Service is: Stopped." If not, click Stop Service.

Understanding Xgrid Service

Xgrid service handles the transferring of computing jobs to the grid and returns the results. Xgrid does not calculate anything, know anything about calculating, have content for calculating, or even know that you are calculating anything.

The actual computing job is handled by software (such as perl) that either runs on the network computers, can be installed before running the computing job, or is transferred to the computers using Xgrid.

The primary components of a computational grid perform these functions:

- An agent runs one task at a time per CPU; a dual-processor computer can run two tasks simultaneously.
- A controller queues tasks, distributes those tasks to agents, and handles task reassignment.
- A client submits jobs to the Xgrid controller in the form of multiple tasks. A client can be any computer running Mac OS X v10.4 or Mac OS X Server v10.4.

In principle, the agent, controller, and client can run on the same server, but it is often more efficient to have a dedicated controller node.

Authenticating for the Grid

Xgrid requires controllers to mutually authenticate with both clients and agents. There are three authentication options. By default, Xgrid uses passwords so that nothing happens without authentication. Only authorized users can run jobs on the grid and grid computers only accept jobs that have the proper password. Also, there is only one computer on a grid that has a listening port.

Single Sign-On

Single sign-on (SSO) is the most powerful and flexible form of authentication. It leverages the Open Directory and Kerberos infrastructure in Mac OS X Server v10.4 to manage authentication behind the scenes, without direct user intervention. Each entity has a Kerberos principal, which passes the appropriate ticket to determine identity, which is then checked against the relevant group to determine privileges. Generally, you should use this option if any of the following conditions is true:

- You already have SSO in your environment.
- You have administrative control over all the agents and clients in use.
- Jobs need to run with special privileges (such as for local, network, or SAN file system access).

Password-Based Authentication

You have the option of password authentication when you can't use single sign-on. You may not be able to use SSO if:

- Potential Xgrid clients are not trusted by your SSO domain (or you don't have one).
- You want to use agents that are on the Internet or otherwise outside your control.
- It is an ad hoc grid, without any ability to prearrange a web of trust.

In these situations, your best option is to specify a password. You have two passwords: one for controller-client and one of controller-agent. Ideally these should be different, for security reasons.

Note: It is possible to create hybrid environments, such as with client-controller authentication done with passwords but controller-agent authentication done with SSO (or vice versa).

No Authentication

No authentication should never be used. It creates potential security risks, since anyone can then connect or run a job which can expose sensitive data. This option is appropriate only for testing a private network in a home or lab that is inaccessible from any untrusted computer, or when none of the jobs or the computers contain any sensitive or important information.

Setting Passwords for Xgrid Service

You can use passwords, Kerberos authentication, or no authentication for Xgrid components. You specify password options in Server Admin as part of configuring the agent and controller.

You set up an Xgrid controller in the Server Admin application. You can specify authentication for agents and clients, and the passwords entered in Server Admin for the controller must match those entered for each agent and client.

If you choose to have no authentication, agents can join the grid and clients can submit jobs to the grid without authenticating. This method is not recommended for any grid or job in which security is important, but it can be useful for testing a grid or running practice jobs. Consider these points when establishing passwords for agents and clients:

- **Kerberos authentication (single sign-on).** If you use Kerberos authentication for agents or clients, the server that's the Xgrid controller must already be configured for Kerberos, be in the same realm as the server running the KDC system, and be bound to the Open Directory master. The agent uses the host principal found in the /etc/krb5.keytab file. The controller uses the Xgrid service principal found in the /etc/krb5.keytab file.

- **Agents.** The agent determines the authentication method. The controller must conform to that method and password (if a password is used). When an agent is configured with a standard password (not single sign-on), you must use the same password for agents when you configure the controller. If the agent has specified single sign-on, the appropriate service principal and host principals must be available.
- **Clients.** If your server is the controller for a grid, you must be sure that administrators of Mac OS X and Mac OS X Server clients on your network use the correct authentication for the controller. A client can submit a job to the controller only if its password is the same as that for the controller, or if the appropriate single sign-on (Kerberos) principals are available.

Securely Configuring Xgrid Service

The Xgrid service must be running for your server to control a grid or participate in a grid as an agent. If the Xgrid service is required, you must configure the Xgrid agent and controller. The Xgrid controller and agent are disabled by default.

When configuring the Xgrid agent and controller, it is important to require authentication to protect your network from malicious users. Authentication requires that both the agent and the controller use the same password or authenticate using Kerberos single sign-on. With no authentication, a malicious agent could receive tasks and potentially access sensitive data.

Configuring an Xgrid Agent

The Xgrid agents run the computational tasks of a job. In Mac OS X Server, the agent is turned off by default. When an agent has been turned on and becomes active at startup, it registers with a controller. (An agent can be connected to only one controller at a time.) The controller sends instructions and data to the agent for the controller's jobs. Once it receives instructions from the controller, the agent executes its assigned tasks and sends the results back to the controller.

You use Server Admin to set up the Xgrid agent on your server. In addition to enabling the agent, you can associate the agent with a specific controller or allow it to join any grid, specify when the agent accepts tasks, and set a password that the controller must recognize.

To configure an Xgrid agent:

1. Open Server Admin.
2. In the Computers & Services list, select Xgrid for the server that you are configuring.
3. Click Settings.
4. Click Agent.

5 Click "Enable agent service."

Skip this step if the controller is already checked.

6 Choose the options you want for the agent.

To specify a controller, choose its name in the pop-up menu or enter the controller name.

Note: An agent can find a controller in one of three ways: the administrator specifies one by host name or IP address; the agent binds to the first controller that advertises on mDNS on the local subnet; or a service lookup is performed against the domain name server for _xgrid._tcp._ip.

7 Choose an authentication method from the pop-up menu and enter the password.

Password requires that the agent and controller use the same password.

Kerberos uses the single sign-on authentication for the agent's administrator.

None does not require a password for the agent (not recommended).

8 Click Save to save your changes and restart the service.

Important: If you require authentication, both the agent and the controller must use the same password or must authenticate using Kerberos single sign-on.

Configuring an Xgrid Controller

You use Server Admin to configure an Xgrid controller. Passwords you set for the controller determine whether client computers can submit jobs and whether agents can join a grid.

To configure an Xgrid controller:

1 Open Server Admin.

2 In the Computers & Services list, select Xgrid for the server that you are configuring.

3 Click Settings.

4 Click Controller.

5 Click "Enable controller service."

Skip this step if the agent is already checked.

6 Choose an authentication option for clients from the pop-up menu and enter the password. (Clients submit jobs to the controller.)

Password requires that the agent and controller use the same password.

Kerberos uses the single sign-on authentication for the agent's administrator.

None does not require a password for the agent (not recommended).

Important: If you require authentication, both the agent and the controller must use the same password or must authenticate using Kerberos single sign-on.

7 Choose an authentication option for agents from the pop-up menu and enter the password.

Password requires that the agent and controller use the same password.

Kerberos uses the single sign-on authentication for the agent's administrator.

Any allows the controller to use either the agent's specified password or the single sign-on password.

None does not require a password for the agent (not recommended).

8 Click Save to save your changes and restart the service.

Validating System Integrity

21

Monitoring events and logs can help to protect the integrity of your computer.

Using auditing and logging tools to monitor your computer can help you secure your computer. By reviewing these audits and log files, you can stop login attempts from unauthorized users or computers and further protect your configuration settings. This chapter also discusses antivirus tools, which detect unwanted viruses.

Using Activity Analysis Tools

Mac OS X includes several command-line tools that you can use to analyze computer activity.

Depending on the tools' configurations and your computer's activity, running these tools can use a lot of disk space. Additionally, these tools are only effective when other users don't have administrator access. Users with administrator access can edit logs generated by the tool and thereby circumvent the tool.

If your computer contains sensitive data, you should consider using both auditing and logging tools. By using both types of tools, you are able to properly research and analyze intrusion attempts and changes in your computer's behavior. You must configure these tools to meet your organization's needs, and then change their logging settings to create only relevant information for reviewing or archiving purposes.

Configuring System Auditing

Auditing is the capture and maintenance of information about security-related events. Auditing helps determine the causes and the methods used for both successful and failed access attempts.

The audit subsystem allows authorized administrators to create, read, and delete audit information. The audit subsystem creates a log of auditable events and allows the administrator to read all audit information from the records in a manner suitable for interpretation. The default location for these files is the /var/audit/ folder.

The audit subsystem is controlled by the audit utility located in the /usr/sbin/ folder. This utility transitions the system in and out of audit operation.

The default configuration of the audit mechanism is controlled by a set of configuration files in the /etc/security/ folder.

If auditing is enabled, the /etc/rc startup script starts the audit daemon at system startup. All the features of the daemon are controlled by the audit utility and audit_control file.

Installing Auditing Tools

The Common Criteria Tools disk image (.dmg) file contains the installer for auditing tools. This disk image file is available from the Common Criteria webpage located at www.apple.com/support/security/commoncriteria/.

After downloading the Common Criteria Tools disk image file, copy it to a removable disk, such as a CD-R disc, FireWire disk, or USB disk.

To install the Common Criteria Tools software:

1. Insert the disk that contains the Common Criteria Tools disk image file and open the file to mount the volume containing the tools Installer.
2. Double-click the CommonCriteriaTools.pkg installer file.
3. Click Continue, then proceed through the installation by following the onscreen instructions.
4. When prompted to authenticate, enter the user name and password of the administrator account.

Enabling Auditing

Modify the hostconfig file to enable auditing.

To turn auditing on:

1. Open Terminal.
2. Enter the following command to edit the /etc/hostconfig file.

```
$ sudo vi /etc/hostconfig
```

3. Add the following entry to the file.

```
AUDIT=-YES-
```

4. Save the file.

Auditing is enabled automatically when the computer starts up.

The following table shows the possible audit settings and what they do.

Parameter	Description
`AUDIT=-YES-`	Enable auditing; ignore failure.
`AUDIT=-NO-`	Disable auditing.
`AUDIT=-FAILSTOP-`	Enable auditing; processes may stop if failure occurs.
`AUDIT=-FAILHALT-`	Enable auditing; the system will be halted if failure occurs.

If the `AUDIT` entry is missing from the /etc/hostconfig file, then auditing is turned off. A failure is any occurrence that prevents audit events from being logged.

The audit subsystem generates warnings when relevant events, such as storage space exhaustion and errors in operation are recognized during audit startup or log rotation. These warnings are communicated to the audit_warn script, which can then communicate these events to the authorized administrator.

Setting Audit Mechanisms

The system startup scripts attempt to configure auditing early in the system startup process. After auditing is enabled, the settings for the audit mechanism are set with the /etc/security/audit_control configuration file.

Files containing audit settings can be edited with any text editor. Terminal can be used in conjunction with `vi` or `emacs` text editor tools. For more information about using text editors with Terminal, see the `vi` or `emacs` man page.

Audit flags are defined in terms of audit classes. Audit flags can be for the whole system, or specific flags can be used for a particular user. Audit flags can include or exclude classes of events from the audit record stream based on the outcome of the event. For example, the outcome could be success, failure or both.

When a user logs in, the system-wide audit flags from the audit_control file are combined with the user-specific audit flags (if any) from the audit_user file, and together establish the preselection mask for the user. The preselection mask determines which events will generate audit records for the given user. If the preselection mask is changed, you should restart the computer to ensure that all components are producing audit events consistently.

Using the audit Tool

Auditing is managed by the `audit` tool. The `audit` tool uses this syntax:

```
$ audit [-nst] [file]
```

The `audit` tool controls the state of the auditing subsystem. The optional file operand specifies the location of the audit_control input file. The default file is /etc/security/audit_control.

You can use the following options with the `audit` tool.

Parameter	Description
`-n`	Forces the audit system to close the existing audit log file and rotate to a new log file in a location specified in the audit control file.
`-s`	Specifies that the audit system should [re]start and reread its configuration from the audit control file. A new log file is created.
`-t`	Specifies that the audit system should terminate. Log files are closed and renamed to indicate the time of the shutdown.

For more information, see the `audit` man page.

Using the auditreduce Tool

The `auditreduce` tool allows you to select events that have been logged in the audit records. Matching audit records are printed to the standard output in their raw binary form. If no filename is specified, the standard input is used by default.

The `auditreduce` tool follows this syntax:

```
$ auditreduce [-A] [-a YYYYMMDD[HH[MM[SS]]]] [-b YYYYMMDD[HH[MM[SS]]]]
     [-c flags] [-d YYYYMMDD] [-e euid] [-f egid] [-g rgid] [-r ruid]
     [-u auid] [-j id] [-m event] [-o object=value] [file ...]
```

For more information, see the `auditreduce` man pages.

Parameter	Description
`-A`	Selects all records.
`-a`	`YYYYMMDD [HH[MM[SS]]]` Selects records that occurred after or on the given date and time.
`-b`	`YYYYMMDD [HH[MM[SS]]]` Selects records that occurred before the given date and time.
`-c`	`flags` Selects records matching the given audit classes specified as a comma-separated list of audit flags.
`-d`	`YYYYMMDD` Selects records that occurred on a given date. Cannot be used with `-a` or `-b` option flags.
`-e`	`euid` Selects records with the given effective user.
`-f`	`egid` Selects records with the given effective group.
`-g`	`gid` Selects records with the given real group.

Parameter	Description
`-r`	`ruid` Selects records with the given real user.
`-u`	`auid` Selects records with the given audit ID.
`-j`	`id` Selects records having a subject token with matching ID.
`-m`	`event` Selects records with the given event name or number.
`-o`	`object` = value `file` = Selects records containing the given path name. `file` ="/usr" matches paths starting with usr. `file` ="~/usr" matches paths not starting with usr. `msgqid` = Selects records containing the given message queue id. `pid` = Selects records containing the given process id. `semid` = Selects records containing the given semaphore id. `shmid` = Selects records containing the given shared memory id.

To select all records associated with effective user ID root from the audit log /var/audit/20031016184719.20031017122634:

```
$ auditreduce -e root /var/audit/20031016184719.20031017122634
```

To select all setlogin events from that log:

```
$ auditreduce -m AUE_SETLOGIN /var/audit/20031016184719.20031017122634:
```

Using the praudit Tool

The `praudit` tool prints the contents of the audit records. The audit records are displayed in standard output (stdout). If no filename is specified, standard input (stdin) is used.

The `praudit` tool uses this syntax:

```
$ praudit [options] audit-trail-file [....]
```

You can use `praudit` with the following options:

Parameter	Description
`-l`	Prints the entire record in the same line. If this option is not specified, every token is displayed in a different line.
`-r`	Prints records in their raw format. This option is separate from `-s`.
`-s`	Prints the tokens in their "short" form. Short ASCII representations for record and event type are displayed. This option is separate from `-r`.
`del`	Specifies the delimiter. The default delimiter is the comma.

If raw or short form are not specified, tokens are printed in their long form. That is, events are displayed according to their descriptions given in audit_event; UIDs and GIDs are expanded to their actual ASCII representation, date and time is displayed in standard date format, and so on.

For more information, see the `praudit` man page.

Deleting Audit Records

You can clear the audit trail by deleting audit files using the command line.

> *WARNING:* You should not delete the currently active audit log.

To delete an audit file:

```
$ sudo rm /var/audit/20031016184719.20031017122634
```

Audit Control Files

There are several text files the audit system uses to control auditing and write audit records. The default location for these files is the /etc/security/ folder.

- audit_class—The audit_class file contains descriptions of the auditable event classes on the system. Each auditable event is a member of an event class. Each line maps an audit event mask (bitmap) to a class and a description.
- audit_control—The audit_control file contains several audit system parameters. Each line of this file is of the form `parameter:value`. Audit flags are a comma-delimited list of audit classes as defined in the audit_class file. Event classes can be preceded by a prefix that changes their interpretation.
- audit_event—The audit_event file contains descriptions of the auditable events on the system. Each line maps an audit event number to a name, a description, and a class. Each event class should have a corresponding entry in the audit_class file.
- audit_user—The audit_user file specifies which audit event classes are to be audited for the given users. If specified, these flags are combined with the system-wide audit flags in the audit_control file to determine which classes of events to audit for that user. These settings take effect when the user logs in. Each line maps a user name to a list of classes that should be audited and a list of classes that should not be audited.
- audit_warn—The audit_warn file runs when auditd generates warning messages. The default audit_warn is a script whose first parameter is the type of warning; the script appends its arguments to /etc/security/audit_messages. Administrators can replace this script with a more comprehensive one that takes different actions based on the type of warning. For example, a low-space warning could result in an email message being sent to the administrator.

For more information about editing audit control files, see the common criteria administration guide at www.apple.com/support/security.

Managing Audit Log Files

If auditing is enabled, the auditing subsystem adds records of auditable events to an audit log file. The name of an audit log file consists of the date and time it was created, followed by a period, and the date and time it was terminated. For example:

20040322183133.20040322184443.

This log was created on March 22nd 2004 at 18:31:33 and was terminated on March 22nd 2004 at 18:44:43.

The audit subsystem appends records to only one audit log file at any given time. The currently active file has a suffix ".not_terminated" instead of a date and time. Audit log files are stored in the folders specified in the audit_control file. The audit subsystem creates an audit log file in the first folder specified.

When less than the minfree amount of disk space is available on the volume containing the audit log file, the audit subsystem:

- Issues an audit_warn soft warning
- Terminates the current audit log file
- Creates a new audit log file in the next specified folder

Once all folders specified have exceeded this minfree limit, auditing resumes in the first folder again. However, if that folder is full, an auditing subsystem failure can occur.

You can also choose to terminate the current audit log file and create a new one manually using the audit utility. This action is commonly referred to as "rotating the audit logs."

Use `audit -n` to rotate the current log file. Use `audit -s` to force the audit subsystem to reload its settings from the audit_control file (this also rotates the current log file).

Configuring Log Files

Logging is the recording of various events, including changes to service status, processes, and operating system components. Some of these events are security related, while others are information messages about your computer's activity. If an unexpected error occurs, you can analyze logs to help determine the cause of the error. For example, the logs might explain why a software update can't be installed, or why you can't authenticate.

Logging tools can be useful if you have multiple users who can access the `sudo` command. You can view logs to see what users did using the `sudo` command. Some `sudo` commands perform additional actions that are not logged. You should restrict the `sudo` commands that individual users are allowed to use. For more information, see "Restricting sudo Usage" on page 71.

Use Console to view and maintain log files. Console is located in the /Applications/Utilities/ folder. Upon starting, the Console window shows the `console.log` file. Click Logs to display a pane that shows other log files on the system in a tree view. The tree includes folders for services, such as web and email server software.

In Mac OS X Server, log files are handled by either the BSD subsystem or a specific application. The BSD subsystem handles most of the important system logging, while some applications will handle their own logging. Like other BSD systems, Mac OS X Server uses a background process called `syslogd` to handle logging. A primary decision to make when configuring `syslogd` is whether to use local or remote logging. In local logging, log messages are stored on the hard disk. In remote logging, log messages are transferred over the network to a dedicated log server that stores them. Using remote logging is strongly recommended for any system.

Configuring the syslogd Daemon

The configuration file for the system logging process, `syslogd`, is /etc/syslog.conf. A manual for configuration of this file is available by issuing the command `man syslog.conf` in a Terminal window. Each line within /etc/syslog.conf consists of text containing three types of data: a facility, a priority, and an action. Facilities are categories of log messages. The standard facilities include mail, news, user, and kern (kernel). Priorities deal with the urgency of the message. In order from least to most critical, they are: debug, info, notice, warning, err, crit, alert, and emerg. The priority of the log message is set by the application sending it, not by `syslogd`. Finally, the action specifies what to do with a log message of a specific facility and priority. Messages can be sent to files, named pipes, devices, or a remote host.

The following example specifies that for any log messages in the category "mail," with a priority of "emerg" or higher, the message is written to the /var/log/mail.log file:

```
mail.emerg /var/log/mail.log
```

The facility and priority are separated by only a period, and these are separated from the action by one or more tabs. Wildcards ("*") can also be used in the configuration file.

The following example logs all messages of any facility or priority to the file /var/log/all.log:

```
*.* /var/log/all.log
```

Local System Logging

The default configuration in /etc/syslog.conf is configured for local logging in the /var/log folder. The system is set to rotate log files using a cron job at the time intervals specified in the /etc/crontab file. Rotation entails compressing the current log file, incrementing the integer in the filename of compressed log files, and creating a new log file for new messages.

The following table describes the rotation process after two rotations.

Files before rotation:	Files after first rotation:	File after second rotation:
system.log	system.log	system.log
mail.log	mail.log	mail.log
	mail.log.1.gz	mail.log.1.gz
	system.log.1.gz	system.log.1.gz
		mail.log.2.gz
		system.log.2.gz

The log files are rotated by a cron job, and the rotation only occurs if the system is on when the job is scheduled. By default, the log rotation tasks are scheduled for very early in the morning (for example, 4:30 a.m. on Saturday) to be as unobtrusive as possible to the user. If the system will not be on at this time, adjust the settings in /etc/crontab.

For information about editing the /etc/crontab file, enter `man 5 crontab` in a Terminal window.

The following line shows the default for running the weekly log rotation script, which is configured for 4:15 a.m. on the last day of the week, Saturday (Sunday is 0). An asterisk denotes "any," so a line of all asterisks would execute every minute.

```
             DayOf       DayOf
#Minute Hour Month Month Week User Command
     15 4    *     *     6    root periodic weekly
```

The following line would change the time to 12:15 p.m. on Tuesday, when the system is much more likely to be on:

```
             DayOf       DayOf
#Minute Hour Month Month Week User Command
     15 12   *     *     2    root periodic weekly
```

Remote System Logging

Using remote logging in addition to local logging is strongly recommended for any server system, because local logs can easily be altered if the system is compromised. Several security issues must also be considered when making the decision to use remote logging. First, the syslog process sends log messages in the clear, which could expose sensitive information. Second, too many log messages fill storage space on the logging system, rendering further logging impossible. Third, log files can indicate suspicious activity only if a baseline of normal activity has been established, and if they are regularly monitored for such activity. If these security issues outweigh the security benefit of remote logging for the network being configured, then remote logging should not be used.

The following instructions assume a remote log server has been configured on the network.

To enable remote logging:

1. Open /etc/syslog.conf as root.
2. Add the following line to the top of the file, replacing `your.log.server` with the actual name or IP address of the log server. Make sure to keep all other lines intact:

   ```
   *.* @your.log.server
   ```

3. Exit, saving changes.
4. Send a hangup signal to `syslogd` to make it reload the configuration file:

   ```
   $ sudo killall -HUP syslogd
   ```

Viewing Logs in Server Admin

Server Admin provides logging for some services enabled on your server. A filter feature allows you to search through the log for specific information.

To view logs in Server Admin:

1. Open Server Admin.
2. In the Computers & Services pane, select the service under the server you are viewing logs for.
3. Click Logs.

   Some services have multiple logs associated with them.

About File Integrity Checking Tools

File integrity tools help protect your computer by detecting and logging all changes to file system objects, such as files and folders. Some file integrity tools can also detect changes to your local directory domain, and to any kernel modules. Depending on the file integrity tool you choose, you can also use advanced features, such as the ability to reverse individual file system changes, or to receive highly detailed logs in a variety of formats.

File integrity tools are generally hosted on a server that you securely connect your client to. The server retrieves logs from all clients, and stores baseline configuration databases and current configuration data.

For more information about checksums and file hashing, see "Verifying the Integrity of Software" on page 45.

About Antivirus Tools

Installing antivirus tools helps prevent infection of your computer by viruses, and help prevent your computer from becoming a host for spreading viruses to other computers. These tools quickly identify suspicious content and compare them to known malicious content.

Mac OS X Server uses ClamAV to scan for viruses. If a suspected virus is found, you can choose to deal with it several ways. For information about running email filters, see "Enabling Virus Filtering" on page 233. The virus definitions are kept up to date (if enabled) via the Internet using a process called "freshclam." For more information about ClamAV, see www.clamav.net.

In addition to using antivirus tools, you should develop computer usage habits that are not prone to virus infection. For example, don't download or open content you didn't specifically request, and never open a file sent to you by someone you don't know. For more information about securely using email, see "Enabling Mail Filtering" on page 232.

When you use antivirus tools, make sure you have the latest virus definition files. The protection provided by your antivirus tool depends on the quality of your virus definition files. If your antivirus tool supports it, enable automatic downloading of virus definitions.

For a list of antivirus tools, see the *Macintosh Products Guide* at guide.apple.com.

Understanding Passwords and Authentication

A

Appendix

This appendix explains the different types of passwords and how they are used to authenticate users.

Passwords are a common method for authenticating with another computer. There are several types of services that use passwords to verify the identity of users.

Understanding Password Types

Each user account has a password type that determines how the user account is authenticated. In a local directory domain, the standard password type is shadow password. On a server upgraded from Mac OS X Server version 10.3, user accounts in the local directory domain can also have an Open Directory password type.

For user accounts in the LDAP directory of Mac OS X Server, the standard password type is Open Directory. User accounts in the LDAP directory can also have a password type of crypt password.

Authentication and Authorization

Services, such as the login window and Apple file service, request user authentication from Open Directory. Authentication is part of the process by which a service determines whether or not to grant a user access to a resource. Usually this process also requires authorization. Authentication proves a user's identity, and authorization determines what the authenticated user is allowed to do. A user typically authenticates by providing a valid user name and password. A service can then authorize the authenticated user to access specific resources. For example, file service authorizes full access to folders and files that an authenticated user owns.

You experience authentication and authorization when you use a credit card. The merchant authenticates you by comparing your signature on the sales slip to the signature on your credit card. Then the merchant submits your authorized credit card account number to the bank, which authorizes payment based on your account balance and credit limit.

Open Directory authenticates user accounts and service access control lists (SACLs) authorize use of services. If Open Directory authenticates you, the SACL for login window determines whether you can log in, the SACL for AFP service determines whether you can connect for file service, and so on. Some services also determine whether a user is authorized to access particular resources. This authorization may require retrieving additional user account information from the directory domain. For example, AFP service needs the user ID and group membership information to determine which folders and files the user is authorized to read or write.

Open Directory Passwords

When a user's account has a password type of Open Directory, the user can be authenticated by Kerberos or the Open Directory Password Server. Kerberos is a network authentication system that uses credentials issued by a trusted server. Open Directory Password Server supports the traditional password authentication methods that some clients of network services require. (Kerberos isn't available on some Open Directory servers, such as an upgraded server with a shared NetInfo directory instead of an LDAP directory.)

Neither Kerberos nor Open Directory Password Server stores the password in the user's account. Both Kerberos and Open Directory Password Server store passwords in secure databases apart from the directory domain and never allow passwords to be read. Passwords can only be set and verified. Malicious users might attempt to log in over the network hoping to gain access to Kerberos and Open Directory Password Server. The Open Directory logs can alert you to unsuccessful login attempts.

User accounts in the following directory domains can have Open Directory passwords:

- The LDAP directory of Mac OS X Server
- The local directory domain of Mac OS X Server upgraded from v10.2–10.3
- A shared NetInfo directory of a server upgraded from or still using Mac OS X Server v10.2

Note: Open Directory passwords can't be used to log in to Mac OS X version 10.1 or earlier. Accounts of users who must log in using the login window of Mac OS X v10.1 or earlier must be configured to use crypt passwords. The password type doesn't matter for other services. For example, a user of Mac OS X v10.1 could authenticate for Apple file service with an Open Directory password.

Shadow Passwords

Shadow passwords support the same traditional authentication methods as Open Directory Password Server. These authentication methods are used to send shadow passwords over the network in a scrambled form, or hash.

A shadow password is stored as several hashes in a file on the same computer as the directory domain where the user account resides. Because the password is not stored in the user account, the password is not easy to capture over the network. Each user's shadow password is stored in a different file, called a shadow password file, and these files are protected so they can be read only by the root user account.

Only user accounts that are stored in a computer's local directory can have a shadow password. User accounts that are stored in a shared directory can't have a shadow password.

Shadow passwords also provide cached authentication for mobile user accounts. See the user management guide for information about mobile user accounts.

Crypt Passwords

A crypt password is stored an encrypted value, or hash, in the user account. This strategy, historically called basic authentication, is most compatible with software that must access user records directly. For example, Mac OS X version 10.1 and earlier expect to find a crypt password stored in the user account.

Crypt authentication supports a maximum password length of only eight bytes (eight ASCII characters). If a longer password is entered in a user account, only the first eight bytes are used for crypt password validation. Shadow passwords and Open Directory passwords are not subject to this length limit.

For secure transmission of passwords over a network, crypt supports the DHX authentication method.

Crypt passwords are not secure. They should be used only for user accounts that must be compatible with UNIX clients that require them, or with Mac OS X v10.1 clients. Because they're stored in user accounts, they're too accessible and therefore subject to offline attack. Although stored in an encoded form, they're relatively easy to decode.

Crypt passwords are not stored in clear text. They are concealed and made unreadable by encryption. A crypt password is encrypted by feeding the clear text password along with a random number to a mathematical function, known as a one-way hash function. A one-way hash function always generates the same encrypted value from particular input, but it cannot be used to recreate the original password from the encrypted output it generates.

To validate a password using the encrypted value, Mac OS X applies the function to the password entered by the user and compares it with the value stored in the user account or shadow file. If the values match, the password is considered valid.

Offline Attacks on Passwords

Because crypt passwords are stored directly in user accounts, they are potentially subject to cracking. User accounts in a shared directory domain are accessible on the network. Anyone on the network who has Workgroup Manager or knows how to use command-line tools can read the contents of user accounts, including the passwords stored in them. Note that Open Directory passwords and shadow passwords aren't stored in user accounts, so these passwords can't be read from directory domains. A malicious attacker, could use Workgroup Manager or UNIX commands to copy user records to a file. The cracker can transport this file to a system and use various techniques to decode crypt passwords stored in the user records. After decoding a crypt password, the cracker can log in unnoticed with a legitimate user name and crypt password.

This form of attack is known as an offline attack, since it does not require successive login attempts to gain access to a system. A very effective way to thwart password cracking is to use good passwords. A password should contain letters, numbers, and symbols in combinations that won't be easily guessed by unauthorized users. Passwords should not consist of actual words. Good passwords might include digits and symbols (such as # or $), or they might consist of the first letter of all the words in a particular phrase. Use both uppercase and lowercase letters.

Important: Shadow passwords and Open Directory passwords are far less susceptible to offline attacks because they are not stored in user records. Shadow passwords are stored in separate files that can be read only by someone who knows the password of the root user. Open Directory passwords are stored securely in the Kerberos KDC and in the Open Directory Password Server database. A user's Open Directory password can't be read by other users, not even by a user with administrator rights for Open Directory authentication. (This administrator can change only Open Directory passwords and password policies.)

Password Guidelines

Many applications and services require that you create passwords to authenticate. Mac OS X includes applications that help create complex passwords (Password Assistant), and securely store your passwords (Keychain Access).

Creating Complex Passwords

Use the following tips to create complex passwords:

- Use a mixture of alphabetic (upper and lower case), numeric, and special characters (such as ! and @).
- Don't use words or combinations of words found in a dictionary of any language. Also, don't use names or anything else that is intelligible.

- Create a password of at least twelve characters. Longer passwords are generally more secure than shorter passwords.
- Create as random a password as possible.

You can use Password Assistant to verify the complexity of your password. For more information, see "Using Password Assistant" on page 73.

Using an Algorithm to Create a Complex Password

Consider creating an algorithm to make a complex (but memorable) password. Using an algorithm can increase the randomness of your password. Additionally, instead of having to remember a complex password, you must remember only the algorithm.

The following example shows one possible algorithm for creating a complex password. Instead of using this algorithm, create your own or modify this one.

The following is an algorithm for creating a complex password:

1. Choose your favorite phrase or saying.

   In this example, we'll use:

   Four score and seven years ago our fathers brought forth

   Ideally you should choose a phrase of at least eight words.

2. Reduce your favorite phrase to an acronym by keeping only the first letter of each word.

   The sample phrase becomes:

   Fsasyaofbf

3. Replace a letter with a number.

   If we replace "F" and the last "f" (from "four" and "forth") with "4", and "s" (from "seven") with "7," the sample phrase becomes:

   4sa7yaofb4

4. Add special characters.

   If we add "$" after "4," and "&" after "7," the sample phrase becomes:

   4$sa7&yaofb4$

5. Make some letters uppercase.

   If we convert all vowels to uppercase, the sample phrase becomes:

   4$sA7&yAOfb4$

Safely Storing Your Password

If you store your password or the algorithm used to make your password in a safe place, you'll be able to create more complex passwords without the fear of being unable to recover forgotten passwords. When storing passwords, make sure your storage location is safe, unknown, and inaccessible to intruders. Consider storing your passwords in a sealed envelope within a locked container. Alternatively, you can store your passwords in your wallet. By keeping your passwords in your wallet, you keep passwords in a safe location that is also convenient.

Don't store your password anywhere near your computer.

When writing down your password, take the following precautions:

- Don't identify the password as being a password.
- Don't include account information on the same piece of paper.
- Add some false characters or misinformation to the written password in a way that you remember. Make the written password different from the real password.
- Never record a password online, and never send a password to another person through email.

You can use Keychain Access to store your more complex, longer passwords. You'll still need a password to unlock Keychain Access so that you can view and use these passwords. Because Keychain Access requires that you authenticate to unlock keychains, it is both convenient for you and inaccessible to intruders. Store the Keychain Access password in a safe location. For more information, see "Storing Credentials" on page 75.

Password Maintenance

After you create a good password and store it in a safe location, do the following to make sure your password remains secure:

- Never tell anyone your password. If you tell someone your password, immediately change your password.
- Change your password frequently, and whenever you think your password might be compromised. If your account is compromised, notify authorities and close the account.
- Be aware of when trusted applications ask for your password. Malicious applications can mimic a trusted application and ask you for your password when you're not expecting it.
- Don't reuse the same password for multiple accounts. Otherwise an intruder who compromises your password can use the password for all of those accounts.
- Don't enter password-related hints in "password hint" fields. By providing a hint, you compromise the integrity of your password.

- Don't access your account on public computers or other computers that you don't trust. Malicious computers can record your keystrokes.
- Don't enter your password in front of other people.

Authentication Services

Open Directory offers a variety of options for authenticating users whose accounts are stored in directory domains on Mac OS X Server, including Kerberos and the traditional authentication methods that network services require.

Open Directory can authenticate users by:

- Using Kerberos authentication for single sign-on.
- Using traditional authentication methods and a password stored securely in the Open Directory Password Server database.
- Using traditional authentication methods and a shadow password stored in a secure shadow password file for each user.
- Using a crypt password stored directly in the user's account, for backward compatibility with legacy systems.
- Using a non-Apple LDAP server for LDAP bind authentication.

In addition, Open Directory lets you set up a password policy for all users as well as specific password policies for each user, such as automatic password expiration and minimum password length. (Password policies do not apply to administrators, crypt password authentication, or LDAP bind authentication.)

Determining Which Authentication Option to Use

To authenticate a user, Open Directory must determine which authentication option to use—Kerberos, Open Directory Password Server, shadow password, or crypt password. The user's account contains information that specifies which authentication option to use. This information is called the authentication authority attribute. Open Directory uses the name provided by the user to locate the user's account in the directory domain. Then Open Directory consults the authentication authority attribute in the user's account and learns which authentication option to use.

You can change a user's authentication authority attribute by changing the password type in the Advanced pane of Workgroup Manager, as shown in the following table.

Password Type	Authentication Authority	Attribute in User Record
Open Directory	Open Directory Password Server and/or Kerberos[1]	Either or both: • ;ApplePasswordServer; • ;Kerberosv5;
Shadow password	Password file for each user, readable only by the root user	Either: • ;ShadowHash;[2] • ;ShadowHash;<list of enabled authentication methods>
Crypt password	Encoded password in user record	Either: • ;basic; • no attribute at all

[1] User accounts from Mac OS X Server v10.2 must be reset to include the Kerberos authentication authority attribute.
[2] If the attribute in the user record is ;ShadowHash; without a list of enabled authentication methods, default authentication methods are enabled. The list of default authentication methods is different for Mac OS X Server and Mac OS X.

The authentication authority attribute can specify multiple authentication options. For example, a user account with an Open Directory password type normally has an authentication authority attribute that specifies both Kerberos and Open Directory Password Server.

A user account doesn't have to include an authentication authority attribute at all. If a user's account contains no authentication authority attribute, Mac OS X Server assumes a crypt password is stored in the user's account. For example, user accounts created using Mac OS X version 10.1 and earlier contain a crypt password but not an authentication authority attribute.

Password Policies

Open Directory enforces password policies for users who's password type is Open Directory or Shadow Password. For example, a user's password policy can specify a password expiration interval. If the user is logging in and Open Directory discovers the user's password has expired, the user must replace the expired password. Then Open Directory can authenticate the user.

Password policies are set to prevent unauthorized access. Best password policy practices dictate that you should disable a user account after five failed login attempts, after a month of inactivity, and if the account is temporary on a specified date the account will no longer be needed. Password policies must also require passwords to be a minimum length of twelve characters, contain at least one letter, contain at least one numeral, differ from the account name, differ from the three most recent passwords, and be changed every 90 days.

The password policy for a mobile user account applies when the account is used while disconnected from the network as well as while connected to the network. A mobile user account's password policy is cached for use while offline. For more information about mobile user accounts, see the user management guide.

Password policies do not affect administrator accounts. Administrators are exempt from password policies because they can change the policies at will. In addition, enforcing password policies on administrators could subject them to denial-of-service attacks.

Kerberos and Open Directory Password Server maintain password policies separately. An Open Directory server synchronizes the Kerberos password policy rules with Open Directory Password Server password policy rules.

Single Sign-On Authentication

Mac OS X Server uses Kerberos for single sign-on authentication, which relieves users from entering a user name and password separately for every service. With single sign-on, a user always enters a user name and password in the login window. Thereafter, the user does not have to enter a name and password for Apple file service, mail service, or other services that use Kerberos authentication. To take advantage of the single sign-on feature, users and services must be Kerberized—configured for Kerberos authentication—and use the same Kerberos Key Distribution Center (KDC) server.

User accounts that reside in an LDAP directory of Mac OS X Server and have a password type of Open Directory use the server's built-in KDC. These user accounts are automatically configured for Kerberos and single signon. This server's Kerberized services also use the server's built-in KDC and are automatically configured for single signon. This Mac OS X Server KDC can also authenticate users for services provided by other servers. Having additional servers with Mac OS X Server use the Mac OS X Server KDC requires only minimal configuration.

Kerberos Authentication

Kerberos was developed at MIT to provide secure authentication and communication over open networks like the Internet. It's named for the three-headed dog that guarded the entrance to the underworld of Greek mythology.

Kerberos provides proof of identity for two parties. It enables you to prove who you are to network services you want to use. It also proves to your applications that network services are genuine, not spoofed. Like other authentication systems, Kerberos does not provide authorization. Each network service determines for itself what it will allow you to do based on your proven identity.

Kerberos allows a client and a server to unambiguously identify each other much more securely than the typical challenge-response password authentication methods traditionally deployed. Kerberos also provides a single sign-on environment where users have to authenticate only once a day, week, or period of time, thereby easing authentication loads for the users.

Mac OS X Server and Mac OS X versions 10.3 and 10.4 support Kerberos version 5.

Security Checklist

B Appendix

This appendix contains a checklist of recommended steps required to secure Mac OS X Server.

This appendix contains action item checklists ordered by chapter.

You can customize these checklists to suit your needs. For example, you can mark the completion status of action items in the "Completed?" column. If you deviate from the suggested action item, you can use the "Notes" column to justify or clarify your decision.

Installation Action Items

For details, see Chapter 2, "Installing Mac OS X Server," on page 31.

Action Item	Completed?	Notes
Securely erase the Mac OS X install partition before installation		
Disable the Open Firmware password before installation		
Install Mac OS X using Mac OS Extended disk formatting		
Do not install any unnecessary packages		
Do not transfer confidential information in Server Assistant		
Do not connect to the Internet		
Create administrator accounts with difficult-to-guess names		
Create complex passwords for administrator accounts		
Do not enter a password-related hint, instead enter help desk contact information		

Action Item	Completed?	Notes
Enter correct time settings		
Use an internal Software Update server		
Update system software using verified packages		
Repair disk permissions after installing software or software updates		

Hardware and Core Mac OS X Action Items

For details, see Chapter 3, "Protecting Hardware and Securing Global System Settings," on page 49.

Action Item	Completed?	Notes
Restrict access to rooms that have computers		
Store computers in locked or secure containers when not in use		
Use a password protected screensaver		
Remove Mac OS 9		
When needed, run Mac OS 9 from a CD, DVD, or disc image		
Require an Open Firmware or EFI password		
Create an access warning for the login window		
Create an access warning for the command line		
Do not use fast user switching with nontrusted users or when multiple users access local accounts		

Account Configuration Action Items

For details, see Chapter 4, "Securing Local Server Accounts," on page 63.

Action Item	Completed?	Notes
Create an administrator account and a standard account for each administrator		
Create a standard or a managed account for each nonadministrator		
Set appropriate parental controls for managed accounts		
Restrict the distribution and use of administrator accounts		
Modify the /etc/authorization file to secure directory domain access		
Disable `su`		
Restrict `sudo` users to only being able access required commands		
Set a strong password policy		
Use Password Assistant to help generate complex passwords		
Authenticate using a smart card, token, or biometric device		
Secure the login keychain		
Secure individual keychain items		
Create specialized keychains for different purposes		
Use a portable drive to store keychains		

System Software Action Items

Chapter 5, "Securing System Preferences," describes how to secure system preferences. Every system preference with security-related configuration settings has its own action item checklist.

.Mac Preferences Action Items

For details, see "Securing .Mac Preferences" on page 83.

Action Item	Completed?	Notes
Disable all Sync options		
Disable iDisk Syncing		
Enable Public Folder password protection		
Do not register computers for synchronization		

Accounts Preferences Action Items

For details, see "Securing Accounts Preferences" on page 85.

Action Item	Completed?	Notes
Change initial password for the system administrator account		
Disable automatic login		
Display login window as name and password		
Disable "Show password hints"		
Disable "Enable fast user switching"		
Disable "Show the Restart, Sleep, and Shut Down buttons"		

Appearance Preferences Action Items

For details, see "Securing Appearance Preferences" on page 88.

Action Item	Completed?	Notes
Do not display recent applications		
Do not display recent documents		
Do not display recent servers		

Bluetooth Preferences Action Items

For details, see "Securing Bluetooth Preferences" on page 89.

Action Item	Completed?	Notes
Disable Bluetooth for each user account in System Preferences		
Remove privileges to modify Bluetooth System Preferences		

CDs & DVDs Preferences Actions Items

For details, see "Securing CDs & DVDs Preferences" on page 90.

Action Item	Completed?	Notes
Disable automatic actions for blank CDs for each user account		
Disable automatic actions for blank DVDs for each user account		
Disable automatic actions for music CDs for each user account		
Disable automatic actions for picture CDs for each user account		
Disable automatic actions for video DVDs for each user account		
Remove privileges to modify CDs & DVDs System Preferences		

Classic Preferences Action Items

For details, see "Securing Classic Preferences" on page 90.

Action Item	Completed?	Notes
Disable starting Classic at login		
Do not hide Classic when starting		
Warn before starting Classic		
Show Classic status in the menu bar		
Turn off Classic extensions		
Use the Memory/Versions pane to view all applications running in Mac OS 9		

Dashboard and Exposé Preferences Action Items

For details, see "Securing Dashboard and Exposé Preferences" on page 93

Action Item	Completed?	Notes
Disable Dashboard		

Date & Time Preferences Action Items

For details, see "Securing Date & Time Preferences" on page 94.

Action Item	Completed?	Notes
Set an correct date and time		
Use a secure internal NTP server for automatic date and time setting		

Desktop & Screen Saver Preferences Action Items

For details, see "Securing Desktop & Screen Saver Preferences" on page 95.

Action Item	Completed?	Notes
Set a short inactivity interval for the screen saver		
Set a screen corner to Start Screen Saver for each user account		
Do not set any screen corner to Disable Screen Saver for each user account		
Remove privileges to modify Dashboard and Exposé System Preferences		

Displays Preferences Action Items

For details, see "Securing Displays Preferences" on page 97.

Action Item	Completed?	Notes
Disable display mirroring		

Dock Preferences Action Items

For details, see "Securing Dock Preferences" on page 97.

Action Item	Completed?	Notes
Set the dock to hide when not in use		

Energy Saver Preferences Action Items

For details, see "Securing Energy Saver Preferences" on page 98.

Action Item	Completed?	Notes
Disable sleeping the computer for all power settings		
Enable sleeping the display for all power settings		
Enable sleeping the hard disk for all power settings		
Disable "Wake when the modem detects a ring" for all power settings		
Disable "Wake for Ethernet network administrator access" for power adapter settings		
Disable "Restart automatically after a power failure" for all power settings		
Disable "Restart automatically if the computer freezes" for all power settings		

Keyboard and Mouse Preferences Action Items

For details, see "Securing Keyboard & Mouse Preferences" on page 99.

Action Item	Completed?	Notes
Turn off Bluetooth		

Network Preferences Action Items

For details, see "Securing Network Preferences" on page 100.

Action Item	Completed?	Notes
Disable any unused hardware devices		
Disable IPv6		

Print & Fax Preferences Action Items

For details, see "Securing Print & Fax Preferences" on page 102.

Action Item	Completed?	Notes
Only use printers in secure locations		
Disable printer sharing		
Disable receiving faxes		
Disable sending faxes		

QuickTime Preferences Action Items

For details, see "Securing QuickTime Preferences" on page 103.

Action Item	Completed?	Notes
Disable "Save movies in disk cache"		
Do not install third-party QuickTime software		

Security Preferences Action Items

For details, see "Securing Security Preferences" on page 104.

Action Item	Completed?	Notes
Require a password to wake the computer from sleep or screen saver for each account		

Sharing Preferences Action Items

For details, see "Securing Sharing Preferences" on page 105.

Action Item	Completed?	Notes
Disable Remote Login		
Disable Apple Remote Desktop		
Disable Remote Apple Events		
Rename your computer to a name that does not indicate the purpose of the computer		

Software Update Preferences Action Items

For details, see "Securing Software Update Preferences" on page 106.

Action Item	Completed?	Notes
Disable "Check for updates"		
Disable "Download important updates in the background"		
Manually update using installer packages		
Transfer installer packages from a test-bed computer		
Verify installer packages before installing		

Sound Preferences Action Items

For details, see "Securing Sound Preferences" on page 107.

Action Item	Completed?	Notes
Minimize input volume for the internal microphone		
Minimize input volume for the audio line-in port		

Speech Preferences Action Items

For details, see "Securing Speech Preferences" on page 108.

Action Item	Completed?	Notes
Only enable speech recognition in a secure environment		
Use headphones if you enable text to speech		

Spotlight Preferences Action Items

For details, see "Securing Spotlight Preferences" on page 109.

Action Item	Completed?	Notes
Prevent Spotlight from searching all confidential folders		

Startup Disk Preferences Action Items

For details, see "Securing Startup Disk Preferences" on page 111.

Action Item	Completed?	Notes
Carefully choose the startup volume		

Data Maintenance and Encryption Action Items

For details, see Chapter 6, "Securing Data and Using Encryption," on page 113.

Action Item	Completed?	Notes
Set global permissions using POSIX or ACLs		
Enable FileVault for every user		
Encrypt portable files		
Set global umask by changing the NSUmask settings		
Mandate secure erasing of files		
Mandate secret erasing of partitions		
Mandate securely erasing free space		

Account Policies Action Items

Chapter 7, "Securing Accounts, Share Points, and Network Views," describes how to set up and manage account policies and user accounts, as well as how to configure settings and preferences for clients. Each of these topics with security-related configuration settings has its own action item checklist.

Share Points Action Items

For details, see "Configuring Share Points" on page 130.

Action Item	Completed?	Notes
Enable SSL in Workgroup Manager		
Disable unused share points		
Disable unused sharing protocols		
Restrict share point access		

Network Views Action Items

For details, see "Controlling Network Views" on page 135.

Action Items	Completed	Notes
Configure network views to restrict awareness of servers		

Account Configuration Action Items

For details, see "Securing Accounts" on page 136.

Action Item	Completed?	Notes
Disallow simultaneous login		
Use an Open Directory password instead of a crypt password		
Enter a disk quota		
Use either POP or IMAP for email, not both		
Use POSIX or ACL permissions to determine group account access		
Restrict access to specific groups by assigning computers to a list		
If accounts are stored in a network domain, disable local accounts		
Specify a time interval to update the preferences cache		

Applications Preferences Action Items

For details, see "Managing Applications Preferences" on page 145.

Action Item	Completed?	Notes
Create a list of approved applications which users can open		
Deselect "User can also open all applications on local volumes"		
Deselect "Allow approved applications to launch non-approved applications"		
Deselect "Allow UNIX tools to run"		

Classic Preferences Action Items

For details, see "Managing Classic Preferences" on page 146.

Action Item	Completed?	Notes
Deselect "Start up Classic at login"		
Select "Warn at Classic startup"		
Select "Allow special startup modes"		

Dock Preferences Action Items

For details, see "Managing Dock Preferences" on page 147.

Action Item	Completed?	Notes
Modify the Applications list to include required applications		
Modify the Documents and Folders list to include required documents and folders		
Deselect '"Merge with user's Dock"		
Deselect "My Applications"		
Deselect "Documents"		
Deselect "Network Home"		
Select "Automatically hide and show the Dock"		

Energy Saver Preferences Action Items

For details, see "Managing Energy Saver Preferences" on page 149.

Action Item	Completed?	Notes
Disable sleeping the computer for all power settings		
Deselect "Start up the computer"		

Finder Preferences Action Items

For details, see "Managing Finder Preferences" on page 150.

Action Item	Completed?	Notes
Select "Use normal finder"		
Deselect "Hard Disks"		
Deselect "Removable media (such as CDs)"		
Deselect "Connected Servers"		
Select "Always show file extensions"		
Deselect "Connect to Server"		
Deselect "Go to iDisk"		
Deselect "Go to Folder"		
Deselect "Eject"		
Deselect "Burn Disk"		

Action Item	Completed?	Notes
Deselect "Restart"		
Deselect "Shut Down"		

Internet Preferences Action Items

For details, see "Managing Internet Preferences" on page 152.

Action Item	Completed?	Notes
Enter user email information		
Select POP or IMAP		
Enter approved intranet pages		
Designate a location to download files		

Login Preferences Action Items

For details, see"Managing Login Preferences" on page 155.

Action Item	Completed?	Notes
Deselect "Add network home share point"		
Deselect "User may add and remove additional items"		
Deselect "User may press Shift to keep items from opening"		
Do not allow login or log-out scripts		
Do not allow LoginHook or LogoutHook scripts		
Enter help desk information as the login message		
Display login window as name and password text fields		
Do not allow Restart or Shut Down buttons to show in the Login Window		
Do not allow password hints		
Deselect "Auto Login Client Setting"		
Deselect "Allow users to log in using '>console.'"		

Action Item	Completed?	Notes
Deselect "Enable Fast User Switching"		
Deselect "Log out users after # minutes of activity"		

Media Access Preferences Action Items

For details, see "Managing Media Access Preferences" on page 159.

Action Item	Completed?	Notes
Disable unnecessary media		
Deselect Allow for CDs		
Deselect Allow for CD-ROMs		
Deselect Allow for DVDs		
Deselect Allow for Recordable Disks		
Deselect Allow for Internal Disks		
Deselect Allow for External Disks		
Select "Eject all removable media at logout"		

Mobility Preferences Action Items

For details, see "Managing Mobility Preferences" on page 161.

Action Item	Completed?	Notes
Do not use mobile account on insecure or infrequently accessed computers		
Use FileVault on every computer with portable home folders		
Deselect "Synchronize account for offline use"		

Network Preferences Action Items

For details, see "Managing Network Preferences" on page 163.

Action Item	Completed?	Notes
Use your organization-controlled proxy servers		
Bypass trusted hosts and domains		
Deselect "Use Passive FTP Mode (PASV)"		

Printing Preferences Action Items

For details, see "Managing Printing Preferences" on page 165.

Action Item	Completed?	Notes
Reduce access to printers		
Deselect "Allow user to modify the printer list"		
Deselect "Allow printers that connect directly to user's computer"		
If selecting "Allow printers that connect directly to user's computer", then select "Require an administrator password"		
Select a printer and select "Require an administrator password"		

Software Update Preferences Action Items

For details, see "Managing Software Update Preferences" on page 167.

Action Item	Completed?	Notes
Designate an internal server to control software updates		

System Preferences Preferences Action Items

For details, see "Managing System Preferences Preferences" on page 168.

Action Item	Completed?	Notes
Select Appearance to appear in the System Preferences preferences		
Select Dashboard & Exposé to appear in the System Preferences preferences		
Select Displays to appear in the System Preferences preferences		
Select Dock to appear in the System Preferences preferences		
Select Keyboard & Mouse to appear in the System Preferences preferences		
Select Security to appear in the System Preferences preferences		

Action Item	Completed?	Notes
Select Universal to appear in the System Preferences preferences		
Disable widgets for network managed users		

Universal Access Preferences Action Items

For details, see "Managing Universal Access Preferences" on page 169.

Action Item	Completed?	Notes
Deselect "Turn on Zoom"		
Set Sticky Keys to Off		
Deselect "Show pressed keys on screen"		

Certificates Action Items

For details, see Chapter 8, "Managing Certificates," on page 171.

Action Item	Completed?	Notes
Obtain certificates to use with SSL-enabled services		
Create a CA to issue certificates		
Create an SSL certificate for distribution		
Create the files and folders needed by SSL		
Export certificate to client computers		

General Protocols and Service Access Action Items

For details, see Chapter 9, "Setting General Protocols and Access to Services," on page 183.

Action Item	Completed?	Notes
Disable NTP		
Disable SNMP		
Enable SSH		
Do not use "server" or your name to identify the server		
Set an correct date and time		

Action Item	Completed?	Notes
Use a secure internal NTP server for automatic date and time setting		
Use Certificate Manager to create, use, and maintain identities for SSL-enabled services		
Use SACL to restrict access to AFP, FTP, and Windows file services		

Remote Access Services Action Items

For details, see Chapter 10, "Securing Remote Access Services," on page 191.

Action Item	Completed?	Notes
Disable root login using SSH		
Modify the /private/etc/sshd_config file to further SSH		
Generate identity key pairs for login authentication		
Configure access for using SSH through Server Admin using SACLs		
Use SFTP instead of FTP		
Disable VPN services		
If using VPN services, enable either or both L2TP and PPTP transport protocols		
To use SecurID authentication, edit the VPN configuration file manually		
Configure an access warning banner		
Disable Apple Remote Desktop		
Encrypt all Observe and Control traffic by setting "Encrypt all network data"		
Encrypt network data during file copy and package installation by setting "Encrypt transfers when using Install Packages"		
Disable Remote Apple Events		

Network and Host Access Services Action Items

Chapter 11, "Securing Network and Host Access Services," describes configuration information to secure your network services. Several services are provided to maintain your network. Each of these services with security-related configuration settings has its own action item checklist.

IPv6 Protocol Action Items

For details, see "Using IPv6 Protocol" on page 205.

Action Item	Completed?	Notes
Enable IPv6		
Configure IPv6 manually or automatically		

DHCP Service Action Items

For details, see "Securing DHCP Service" on page 206.

Action Item	Completed?	Notes
Disable the DHCP service if not required		
If using DHCP, disable DNS, LDAP, and WINS		
Assign static IP addresses		

DNS Service Action Items

For details, see "Securing DNS Service" on page 208.

Action Item	Completed?	Notes
Disable the DNS service		
Allow only one system to act as the DNS server		
Allow recursive queries and zone transfers only from trusted clients, not from external networks.		
Update and audit DNS regularly		
Specify which IP addresses are allowed to request zone transfers		
Configure BIND to respond with something other than the current version		
Limit or disable DNS recursion		

Firewall Service Action Items

For details, see "Securing Firewall Service" on page 213.

Action Item	Completed?	Notes
Create IP address groups		
Configure firewall rules for groups and services		
Configure advanced rules for groups and services		
Enable stealth mode		
Set up logging		

NAT Service Action Items

For details, see "Securing NAT Service" on page 220.

Action Item	Completed?	Notes
Disable NAT service if not required		
Configure NAT service		
If necessary, forward incoming traffic to an IP address		

Bonjour Service Action Items

For details, see "Securing Bonjour Service" on page 222.

Action Item	Completed?	Notes
Disable any unused services that should not be discovered through Bonjour		

Collaboration Services Action Items

For details, see Chapter 12, "Securing Collaboration Services," on page 223.

Action Item	Completed?	Notes
Disable iChat service		
If using iChat service, designate domain names to use		
Designate a certificate to use		
Monitor communication using iChat service logs		

Mail Service Action Items

For details, see Chapter 13, "Securing Mail Service," on page 227.

Action Item	Completed?	Notes
Turn off support for any protocol that is not required		
Use different systems for providing outgoing and incoming mail service		
Enable SSL for the mail server		
Create and install a signed mail certificate for outgoing and incoming mail service protocols		
The "require" setting in the SSL support options is recommended		
Configure SMTP authentication requirements to reduce junk mail		
Create a list of approved host servers to relay mail		
Enable junk mail filtering		
Enable virus filtering		
Update the virus database at least twice a day		
Set up a problem report account		
Disable the SMTP banner		

File Services Action Items

Chapter 14, "Securing File Services," describes configuring file sharing services. Each type of file sharing service with security-related configuration settings has its own action item checklist.

Action Item	Completed?	Notes
Disable file sharing services if not required		
Use as few protocols as possible		
Use AFP		
Disable FTP		
Disable NFS		
Disable SMB		

AFP File Sharing Service Action Items

For details, see "Configuring AFP File Sharing Service" on page 237.

Action Item	Completed?	Notes
Disable Bonjour registration		
Disable browsing with AppleTalk		
Disable Guest access		
Enable secure connections		
Disable administrator to masquerade as another user		
Enter "1" for Guest Connections		
Enable access log		
Set frequency of archiving		
Implement settings for idle user		

FTP File Sharing Service Action Items

For details, see "Configuring FTP File Sharing Service" on page 238.

Action Item	Completed?	Notes
If authentication is possible, use SFTP instead of FTP		
Disconnect client after 1 login failure		
Enter an email address specifically set up to handle FTP administration		
Select Kerberos for access authentication		
Allow maximum of 1 authenticated users		
Enable anonymous access and designate the number of anonymous users		
Disable MacBinary and disk image auto-conversion		
Enable "Show Welcome Message"		
Enable "Show Banner Message"		
Log all login attempts		
Set "Authenticated users see:" to FTP root and Share Points		

Action Item	Completed?	Notes
Designate files to share with anonymous users		
Configure the /Library/FTPServer/Configuration/ftpaccess		

NFS File Sharing Service Action Items

For details, see "Configuring NFS File Sharing Service" on page 240.

Action Item	Completed?	Notes
Use NFS only on a secure LAN or when Apple and Windows file sharing systems are unavailable		
Restrict an NFS share point to those systems that require it		
Make the list of export options as restrictive as possible		

SMB/CIFS Action Items

For details, see "Configuring Windows File Sharing Service" on page 241.

Action Item	Completed?	Notes
Do not allow guest access		
Enter the maximum number of clients connections expected		
Set "Log Detail" to at least medium		
Deselect Workgroup Master Browser and Domain Master Browser services		
Turn off WINS registration		

Web Service Action Items

For details, see Chapter 15, "Securing Web Service," on page 243.

Action Item	Completed?	Notes
Disable web service if not required		
Disable web modules if not required		
Disable web options if not required		

Action Item	Completed?	Notes
Create or obtain separate signed certificates for each domain name		
Enable SSL for web service		
If WebDAV is enabled, assign access privileges for the sites and web folders		
Do not allow Web content files and folders to be writable by world		
Configure a realm to allow user access to websites		
Allow users to access weblogs through an SSL enabled site		

Client Configuration Management Services Action Items

For details, see Chapter 16, "Securing Client Configuration Management Services," on page 255.

Action Item	Completed?	Notes
Disable NetBoot and NetBoot Disk Images		
Use Server Admin to view NetBoot clients and the status of NetBoot service		

Directory Services Action Items

For details, see Chapter 17, "Securing Directory Services," on page 259.

Action Item	Completed?	Notes
Configure Open Directory roles		
Configure Kerberos		
Set a server outside of directory domains as Standalone Server		
Enable SSL		
Set global password policies		
Set binding policies		
Set security policies for Open Directory		

Print Service Action Items

For details, see Chapter 18, "Securing Print Service," on page 267.

Action Item	Completed?	Notes
Use Server Admin to manage print queues and configure settings		
Specify a default LPR queue		

Multimedia Services Action Items

For details, see Chapter 19, "Securing Multimedia Services," on page 271.

Action Item	Completed?	Notes
User Server Admin to configure QuickTime streaming service		
Use secure digest authentication to configure client access to streamed media files		

Grid and Cluster Computing Services Action Items

For details, see Chapter 20, "Securing Grid and Cluster Computing Services," on page 277.

Action Item	Completed?	Notes
If possible, use an single sign-on password		
Always require authentication		
Enable Xgrid agent service		
Set a password for your Xgrid		
Enable Xgrid controller service		
Set a password for your Xgrid controller		
Set a password for the server acting as a grid agent		
Set a password for agents to join a grid and clients to submit jobs		

Validating System Integrity Action Items

For details, see Chapter 21, "Validating System Integrity," on page 283.

Action Item	Completed?	Notes
Install and enable auditing tools		
Configure audit settings		
Configure log files		
Configure local system using syslog.conf		
Enable remote system logging		
Install file integrity tools		
Install antivirus tools		

Glossary

This glossary defines terms and spells out abbreviations you may encounter while working with online help or the various reference manuals for Mac OS X Server. References to terms defined elsewhere in the glossary appear in italics.

access control A method of controlling which computers can access a network or network services.

ACE Access Control Entry. An entry within the ACL that controls access rights. See **ACL**.

ACL Access Control List. A list maintained by a system that defines the rights of users and groups to access resources on the system.

administrator A user with server or directory domain administration privileges. Administrators are always members of the predefined "admin" group.

administrator computer A Mac OS X computer onto which you've installed the server administration applications from the Mac OS X Server Admin CD.

AFP Apple Filing Protocol. A client/server protocol used by Apple file service on Macintosh-compatible computers to share files and network services. AFP uses TCP/IP and other protocols to communicate between computers on a network.

authentication The process of proving a user's identity, typically by validating a user name and password. Usually authentication occurs before an authorization process determines the user's level of access to a resource. For example, file service authorizes full access to folders and files that an authenticated user owns.

authentication authority attribute A value that identifies the password validation scheme specified for a user and provides additional information as required.

authorization The process by which a service determines whether it should grant a user access to a resource and how much access the service should allow the user to have. Usually authorization occurs after an authentication process proves the user's identity. For example, file service authorizes full access to folders and files that an authenticated user owns.

BIND Berkeley Internet Name Domain. The program included with Mac OS X Server that implements DNS. The program is also called the name daemon, or named, when the program is running.

binding (n.) A connection between a computer and a directory domain for the purpose of getting identification, authorization, and other administrative data. (v.) The process of making such a connection. See also **trusted binding**.

biometrics A technology that authenticates a person's identity based on unique physiological or behavioral characteristics. Provides an additional factor to authentication. See **two-factor authentication**.

Bonjour A protocol developed by Apple for automatic discovery of computers, devices, and services on IP networks. Formerly called "Rendezvous," this proposed Internet standard protocol is sometimes referred to as "ZeroConf" or "multicast DNS."

BSD Berkeley System Distribution. A version of UNIX on which Mac OS X software is based.

buffer caching Holding data in memory so that it can be accessed more quickly than if it were repeatedly read from disk.

cache A portion of memory or an area on a hard disk that stores frequently accessed data in order to speed up processing times. Read cache holds data in case it's requested by a client; write cache holds data written by a client until it can be stored on disk. See also **buffer caching**, **controller cache**, **disk cache**.

certificate Sometimes called an "identity certificate" or "public key certificate." A file in a specific format (Mac OS X Server uses the x.509 format) that contains the public key half of a public-private keypair, the user's identity information such as name and contact information, and the digital signature or either a *Certificate Authority* (CA) or the key user.

Certificate Authority An authority that issues and manages digital certificates in order to ensure secure transmission of data on a public network. See also **public key infrastructure and certificate.**

cluster A collection of computers interconnected in order to improve reliability, availability, and performance. Clustered computers often run special software to coordinate the computers' activities. See also **computational cluster**.

computational cluster A group of computers or servers that are grouped together to share the processing of a task at a high level of performance. A computational cluster can perform larger tasks than a single computer would be able to complete, and such a grouping of computers (or "nodes") can achieve high performance comparable to a supercomputer.

controller In an Xsan storage area network, short for metadata controller. In RAID systems, controller refers to hardware that manages the reading and writing of data. By segmenting and writing or reading data on multiple drives simultaneously, the RAID controller achieves fast and highly efficient storage and access. See also **metadata controller**.

controller cache A cache that resides within a controller and whose primary purpose is to improve disk performance.

cracker A malicious user who tries to gain unauthorized access to a computer system in order to disrupt computers and networks or steal information. Compare to hacker.

crypt password A type of password that's stored as a hash (using the standard UNIX encryption algorithm) directly in a user record.

daemon A program that runs in the background and provides important system services, such as processing incoming email or handling requests from the network.

decryption The process of retrieving encrypted data using some sort of special knowledge. See also **encryption**.

deploy To place configured computer systems into a specific environment or make them available for use in that environment.

DHCP Dynamic Host Configuration Protocol. A protocol used to dynamically distribute IP addresses to client computers. Each time a client computer starts up, the protocol looks for a DHCP server and then requests an IP address from the DHCP server it finds. The DHCP server checks for an available IP address and sends it to the client computer along with a lease period—the length of time the client computer may use the address.

directory Also known as a folder. A hierarchically organized list of files and/or other directories.

disk cache A cache that resides within a disk. See also **cache**, **controller cache**.

disk image A file that, when opened, creates an icon on a Mac OS X desktop that looks and acts like an actual disk or volume. Using NetBoot, client computers can start up over the network from a server-based disk image that contains system software. Disk image files have a filename extension of either .img or .dmg. The two image formats are similar and are represented with the same icon in the Finder. The .dmg format cannot be used on computers running Mac OS 9.

DNS Domain Name System. A distributed database that maps IP addresses to domain names. A DNS server, also known as a name server, keeps a list of names and the IP addresses associated with each name.

domain Part of the domain name of a computer on the Internet. It does not include the Top Level Domain designator (for example, .com, .net, .us, .uk). Domain name "www.example.com" consists of the subdomain or host name "www," the domain "example," and the top level domain "com."

DoS attack Denial of service attack. An Internet attack that uses thousands of network pings to prevent the legitimate use of a server.

drop box A shared folder with privileges that allow other users to write to, but not read, the folder's contents. Only the owner has full access. Drop boxes should be created only using AFP. When a folder is shared using AFP, the ownership of an item written to the folder is automatically transferred to the owner of the folder, thus giving the owner of a drop box full access to and control over items put into it.

Dynamic Host Configuration Protocol See **DHCP**.

EFI Extensible Firmware Interface. Software that runs automatically when an Intel-based Macintosh first starts up. It determines the computers hardware configuration and starts the system software.

encryption The process of obscuring data, making it unreadable without special knowledge. Usually done for secrecy and confidential communications. See also **decryption**.

Ethernet A common local area networking technology in which data is transmitted in units called packets using protocols such as TCP/IP.

file server A computer that serves files to clients. A file server may be a general-purpose computer that's capable of hosting additional applications or a computer capable only of serving files.

firewall Software that protects the network applications running on your server. IP firewall service, which is part of Mac OS X Server software, scans incoming IP packets and rejects or accepts these packets based on a set of filters you create.

firmware Software that's stored in read-only memory (ROM) on a device and helps in starting up and operating the device. Firmware allows for certain changes to be made to a device without changing the actual hardware of the device.

FTP File Transfer Protocol. A protocol that allows computers to transfer files over a network. FTP clients using any operating system that supports FTP can connect to a file server and download files, depending on their access privileges. Most Internet browsers and a number of freeware applications can be used to access an FTP server.

hacker An individual who enjoys programming, and explores ways to program new features and expand the capabilities of a computer system. See also **cracker**.

hash (noun) A scrambled, or encrypted, form of a password or other text.

host Another name for a server.

host name A unique name for a computer, historically referred to as the UNIX hostname.

HTTP Hypertext Transfer Protocol. The client/server protocol for the World Wide Web. The HTTP protocol provides a way for a web browser to access a web server and request hypermedia documents created using HTML.

ICMP Internet Control Message Protocol. A message control and error-reporting protocol used between host servers and gateways. For example, some Internet software applications use ICMP to send a packet on a round-trip between two hosts to determine round-trip times and discover problems on the network.

image See **disk image**.

IMAP Internet Message Access Protocol. A client-server mail protocol that allows users to store their mail on the mail server rather than download it to the local computer. Mail remains on the server until the user deletes it.

installer package A file package with the filename extension .pkg. An installer package contains resources for installing an application, including the file archive, Read Me and licensing documents, and installer scripts.

IP Internet Protocol. Also known as IPv4. A method used with Transmission Control Protocol (TCP) to send data between computers over a local network or the Internet. IP delivers packets of data, while TCP keeps track of data packets.

IP subnet A portion of an IP network, which may be a physically independent network segment, that shares a network address with other portions of the network and is identified by a subnet number.

IPv4 See **IP**.

IPv6 Internet Protocol version 6. The next-generation communication protocol to replace IP (also known as IPv4). IPv6 allows a greater number of network addresses and can reduce routing loads across the Internet.

JBoss A full-featured Java application server that provides support for Java 2 Platform, Enterprise Edition (J2EE) applications.

KDC Kerberos Key Distribution Center. A trusted server that issues Kerberos tickets.

Kerberos A secure network authentication system. Kerberos uses tickets, which are issued for a specific user, service, and period of time. Once a user is authenticated, it's possible to access additional services without retyping a password (this is called single sign-on) for services that have been configured to take Kerberos tickets. Mac OS X Server uses Kerberos v5.

kernel The part of an operating system that handles memory management, resource allocation, and other low-level services essential to the system.

L2TP Layer Two Tunnelling Protocol. A network transport protocol used for VPN connections. It's essentially a combination of Cisco's L2F and PPTP. L2TP itself isn't an encryption protocol, so it uses IPSec for packet encryption.

LAN Local area network. A network maintained within a facility, as opposed to a WAN (wide area network) that links geographically separated facilities.

LDAP Lightweight Directory Access Protocol. A standard client-server protocol for accessing a directory domain.

managed network The items managed clients are allowed to "see" when they click the Network icon in a Finder window. Administrators control this setting using Workgroup Manager. Also called a "network view."

metadata controller The computer that manages metadata in an Xsan storage area network.

mutual authentication Also known as two-way authentication. A type of authentication in which two parties authenticate with each other. For example, a client or user verifies their identity to a server, and that server confirms its identity to the client or user. Each side has the other's authenticated identity.

NAT Network Address Translation. A method of connecting multiple computers to the Internet (or any other IP network) using one IP address. NAT converts the IP addresses you assign to computers on your private, internal network into one legitimate IP address for Internet communications.

NetBoot server A Mac OS X server on which you've installed NetBoot software and have configured to allow clients to start up from disk images on the server.

Network File System See **NFS**.

network view See **managed network**.

NFS Network File System. A client/server protocol that uses Internet Protocol (IP) to allow remote users to access files as though they were local. NFS exports shared volumes to computers according to IP address, rather than user name and password.

node A processing location. A node can be a computer or some other device, such as a printer. Each node has a unique network address. In Xsan, a node is any computer connected to a storage area network.

NTP Network time protocol. A network protocol used to synchronize the clocks of computers across a network to some time reference clock. NTP is used to ensure that all the computers on a network are reporting the same time.

object class A set of rules that define similar objects in a directory domain by specifying attributes that each object must have and other attributes that each object may have.

offline Refers to data that isn't immediately available, or to devices that are physically connected but not available for use.

Open Directory The Apple directory services architecture, which can access authoritative information about users and network resources from directory domains that use LDAP, NetInfo, or Active Directory protocols; BSD configuration files; and network services.

Open Directory master A server that provides LDAP directory service, Kerberos authentication service, and Open Directory Password Server.

Open Directory password A password that's stored in secure databases on the server and can be authenticated using Open Directory Password Server or Kerberos (if Kerberos is available).

Open Directory Password Server An authentication service that validates passwords using a variety of conventional authentication methods required by the different services of Mac OS X Server. The authentication methods include APOP, CRAM-MD5, DHX, LAN Manager, NTLMv1, NTLMv2, and WebDAV-Digest.

open source A term for the cooperative development of software by the Internet community. The basic principle is to involve as many people as possible in writing and debugging code by publishing the source code and encouraging the formation of a large community of developers who will submit modifications and enhancements.

partition A subdivision of the capacity of a physical or logical disk. Partitions are made up of contiguous blocks on the disk.

PDC Primary domain controller. In Windows networking, a domain controller that has been designated as the primary authentication server for its domain.

permissions Settings that define the kind of access users have to shared items in a file system. You can assign four types of permissions to a share point, folder, or file: read/write, read-only, write-only, and none (no access). See also **privileges**.

phishing An attempt to masquerade as a trusted organization or individual to trick others into divulging confidential information.

PKI Public Key Infrastructure. A mechanism that allows two parties to a data transaction to authenticate each other and use encryption keys and other information in identity certificates to encrypt and decrypt messages they exchange.

POP Post Office Protocol. A protocol for retrieving incoming mail. After a user retrieves POP mail, it's stored on the user's computer and is usually deleted automatically from the mail server.

portable home directory A portable home directory provides a user with both a local and network home folder. The contents of these two home folders, as well as the user's directory and authentication information, can be automatically kept in sync.

POSIX Portable Operating System Interface for UNIX. A family of open system standards based on UNIX, which allows applications to be written to a single target environment in which they can run unchanged on a variety of systems.

print queue An orderly waiting area where print jobs wait until a printer is available. The print service in Mac OS X Server uses print queues on the server to facilitate management.

private key One of two asymmetric keys used in a PKI security system. The private key is not distributed and usually encrypted with a passphrase by the owner. It can digitally sign a message or certificate, claiming authenticity. It can decrypt messages encrypted with the corresponding public key. Finally, it can encrypt messages that can only be decrypted by the private key.

privileges The right to access restricted areas of a system or perform certain tasks (such as management tasks) in the system.

protocol A set of rules that determines how data is sent back and forth between two applications.

proxy server A server that sits between a client application, such as a web browser, and a real server. The proxy server intercepts all requests to the real server to see if it can fulfill the requests itself. If not, it forwards the request to the real server.

public key One of two asymmetric keys used in a PKI security system. The public key is distributed to other communicating parties. It can encrypt messages that can be decrypted only by the holder of the corresponding private key, and it can verify the signature on a message originating from a corresponding private key.

public key certificate See **certificate**.

public key infrastructure A secure method of exchanging data over an unsecure public network, such as the Internet, by using public key cryptography.

QTSS QuickTime Streaming Server. A technology that lets you deliver media over the Internet in real time.

record type A specific category of records, such as users, computers, and mounts. For each record type, a directory domain may contain any number of records.

recursion The process of fully resolving domain names into IP addresses. A nonrecursive DNS query allows referrals to other DNS servers to resolve the address. In general, user applications depend on the DNS server to perform this function, but other DNS servers do not have to perform a recursive query.

rogue computer A computer that is set up by an attacker for the purpose of infiltrating network traffic in an effort to gain unauthorized access to your network environment.

root An account on a system that has no protections or restrictions. System administrators use this account to make changes to the system's configuration.

router A computer networking device that forwards data packets toward their destinations. A router is a special form of gateway which links related network segments. In the small office or home, the term router often means an Internet gateway, often with Network Address Translation (NAT) functions. Although generally correct, the term router more properly refers to a network device with dedicated routing hardware.

RSA Rivest Shamir Adleman algorithm. A public key encryption method that can be used both for encrypting messages and making digital signatures.

SACL Service Access Control List. Lets you specify which users and groups have access to specific services. See **ACL**.

schema The collection of attributes and record types or classes that provide a blueprint for the information in a directory domain.

server A computer that provides services (such as file service, mail service, or web service) to other computers or network devices.

shadow password A password that's stored in a secure file on the server and can be authenticated using a variety of conventional authentication methods required by the different services of Mac OS X Server. The authentication methods include APOP, CRAM-MD5, DHX, LAN Manager, NTLMv1, NTLMv2, and WebDAV-Digest.

share point A folder, hard disk (or hard disk partition), or CD that's accessible over the network. A share point is the point of access at the top level of a group of shared items. Share points can be shared using AFP, Windows SMB, NFS (an "export"), or FTP protocols.

shared secret A value defined at each node of an L2TP VPN connection that serves as the encryption key seed to negotiate authentication and data transport connections.

single sign-on An authentication strategy that relieves users from entering a name and password separately for every network service. Mac OS X Server uses Kerberos to enable single sign-on.

smart card A portable security device that contains a microprocessor. The smart card's microprocessor and its reader use a mutual identification protocol to identify each other before releasing information. The smart card is capable of securely storing passwords, certificates, and keys.

SMB/CIFS Server Message Block/Common Internet File System. A protocol that allows client computers to access files and network services. It can be used over TCP/IP, the Internet, and other network protocols. Windows services use SMB/CIFS to provide access to servers, printers, and other network resources.

SMTP Simple Mail Transfer Protocol. A protocol used to send and transfer mail. Its ability to queue incoming messages is limited, so SMTP usually is used only to send mail, and POP or IMAP is used to receive mail.

SNMP Simple Network Management Protocol. A set of standard protocols used to manage and monitor multiplatform computer network devices.

Spotlight A comprehensive search engine that searches across your documents, images, movies, PDF, email, calendar events, and system preferences. It can find something by its text content, filename, or information associated with it.

SSL Secure Sockets Layer. An Internet protocol that allows you to send encrypted, authenticated information across the Internet. More recent versions of SSL are known as TLS (Transport Level Security).

standalone server A server that provides services on a network but doesn't get directory services from another server or provide directory services to other computers.

static IP address An IP address that's assigned to a computer or device once and is never changed.

streaming Delivery of video or audio data over a network in real time, as a stream of packets instead of a single file download.

subnet A grouping on the same network of client computers that are organized by location (different floors of a building, for example) or by usage (all eighth-grade students, for example). The use of subnets simplifies administration. See also **IP subnet**.

TCP Transmission Control Protocol. A method used along with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. IP takes care of handling the actual delivery of the data, and TCP takes care of keeping track of the individual units of data (called packets) into which a message is divided for efficient routing through the Internet.

ticket, Kerberos A temporary credential that proves a Kerberos client's identity to a service.

trusted binding A mutually authenticated connection between a computer and a directory domain. The computer provides credentials to prove its identity, and the directory domain provides credentials to prove its authenticity.

tunneling A technology that allows one network protocol to send its data using the format of another protocol.

two-factor authentication A process that authenticates through a combination of two independent factors: something you know (such as a password), something you have (such as a smart card), or something you are (such as a biometric factor). This is more secure than authentication that uses only one factor, typically a password.

UDP User Datagram Protocol. A communications method that uses the Internet Protocol (IP) to send a data unit (called a datagram) from one computer to another in a network. Network applications that have very small data units to exchange may use UDP rather than TCP.

VPN Virtual Private Network. A network that uses encryption and other technologies to provide secure communications over a public network, typically the Internet. VPNs are generally cheaper than real private networks using private lines but rely on having the same encryption system at both ends. The encryption may be performed by firewall software or by routers.

WAN Wide area network. A network maintained across geographically separated facilities, as opposed to a LAN (local area network) within a facility. Your WAN interface is usually the one connected to the Internet.

WebDAV Web-based Distributed Authoring and Versioning. A live authoring environment that allows client users to check out webpages, make changes, and then check the pages back in while a site is running.

weblog A webpage that hosts chronologically ordered entries. It functions as an electronic journal or newsletter. Weblog service lets you create weblogs that are owned by individual users or by all members of a group.

workgroup A set of users for whom you define preferences and privileges as a group. Any preferences you define for a group are stored in the group account.

zone transfer The method by which zone data is replicated among authoritative DNS servers. Slave DNS servers request zone transfers from their master servers to acquire their data.

Index

.Mac preferences 83–84
7-pass erase method for data 125–128
35-pass erase method for data 126

A

access
  See also ACLs; LDAP; permissions
  ACEs 46, 117–118
  application 145–146
  Directory Access 136
  Keychain Access 30, 75–79
  remote computers 188
  SACLs 188, 225, 296
  share point 130–134
  startup image configuration 256
  Universal Access 112, 169–170
  weblogs 250–251
  website 249–250
access control entries (ACEs) 46, 117–118
access control lists. *See* ACLs
access file 273
access warnings 59–60
  See also permissions
accounts
  See also user accounts; Workgroup Manager
  administrator 39, 69–71, 136, 263
  computer list 140
  creating secure 64
  group 138–139
  local 140
  mobile 161–163, 297
  overview 30
  preferences 85–87
  types 63–64
ACEs (access control entries) 46, 117–118
ACLs (access control lists)
  keychain services 75
  permissions 45, 46, 113, 117–118, 131
Active Directory 129
activity analysis tools 283–293
addresses. *See* IP addresses; NAT
administrator 46
  accounts for 39, 69–71, 136, 263
  auditing tools 283–289
  directory domain 63, 70, 129, 139
  passwords for 263
  privileges of 29, 46
Advanced Encryption Standard (AES-128) 104
AFP (Apple Filing Protocol) service 130, 133, 237–238
agents, Xgrid 278, 280, 280–281
anonymous access, FTP 238
antivirus tools. *See* virus screening
Apache server 243
Appearance preferences 88
Apple Filing Protocol service. *See* AFP
Apple Remote Desktop (ARD) 105, 203
Apple Software Restore. *See* ASR
AppleTalk 268
applications, user access to 145–146
  See also specific applications
application servers 252–253
ARD. *See* Apple Remote Desktop
ARP (Address Resolution Protocol) spoofing 212
ASR (Apple Software Restore) 38
`asr` tool 38
attributes 133
audit_class file 288
audit_control file 288
audit_event file 288
audit_user file 288
audit_warn file 288
auditing tools 283–289
`auditreduce` tool 286–287
`audit` tool 285–286
authentication
  See also keychain services; passwords
  CHAP 201
  CRAM-MD5 230
  DHX 297
  digest 273
  directory services 129, 265
  EAP 202
  file services 237–238
  Kerberos 278, 279, 301
  LDAP 264

mail 230–231
software updates 106
strengthening methods 72–73
time settings 94
versus authorization 28
Workgroup Manager 129–130
Xgrid 278–280
authentication authority attribute 301–302
authorization 26, 28, 70
See also authentication
automatic login 61

B

banner pages, print 269
Berkeley Internet Name Domain (BIND) 206, 208
Berkeley Software Distribution (BSD) 24, 290
Bill of Materials file 46
BIND (Berkeley Internet Name Domain) 206, 208
binding, trusted 264–265
biometrics-based authentication 74
blogs 250–251
Bluetooth 89
Bonjour browsing service 136, 222
browser security 153
BSD (Berkeley Software Distribution) 24, 290
bundle ID 145

C

cache 140
CA *See* certificates
CDs 33, 52–53, 90
CDSA (Common Data Security Architecture) 24
CERT (Computer Emergency Response Team) 23
Certificate Assistant 175, 178–180
Certificate Authority (CA) 173, 178–181
See also certificates
Certificate Manager 174–177
certificates
certificate authority 173, 178–181
creating 176
deployment to clients 182
iChat server 224
introduction 26
keychain services 121
mail server 228–229
modifying 177–178
Open Directory 262–263
overview 171–173
web service 245–247
Certificate Signing Request (CSR) 171, 174, 175, 176
CHAP (Challenge Handshake Authentication Protocol) 201
chat service 223–225
`chflags` tool 116
`chmod` tool 116, 118
CIFS (Common Internet File System). *See* SMB/CIFS
ClamAV 293
Classic, Mac 51–53, 90–92, 146–147
Classic preferences 146–147
clean installation 33
client computers 30
clients
See also client computers; users
certificate deployment 182
group accounts 138–139
Internet preferences 152–154
Xgrid 278, 280

collaboration services
group accounts 138–139
command-line interface
access warnings 60
Certificate Authority 180
erasing files 127
installing from 37–38
startup security setup 58
Common Criteria Tools 284
Common Data Security Architecture (CDSA) 24
Common Internet File System. *See* SMB/CIFS
Common Security Service Manager (CSSM) 26
Common UNIX Printing System (CUPS) 267, 269
Computer Emergency Response Team (CERT) 23
computer lists 139–140
computer name 186
computers
See also portable computers
client 30
computer lists 139–140
name 186
Console application 290
contacts search policy 136
controllers, Xgrid 278, 281–282
CRAM-MD5 authentication 230
credential storage 75–79
`cron` tool 290
CRYPTOCard KT-1 74
crypt passwords 137, 297, 302
CSR (Certificate Signing Request) 171, 174, 175, 176
CSSM (Common Security Service Manager) 26
CUPS (Common UNIX Printing System) 267, 269
Cyrus 227

D

Dashboard preferences 93–94
databases 129
directory domain 129
Kerberos 264
data security 113–128
Date & Time preferences 94–95, 187
decryption. *See* encryption

denial of service (DoS) attack 212
Desktop preferences 95–97
DHCP (Dynamic Host Configuration Protocol) service 206–208, 264
DHX authentication 297
digest authentication 273
digital signature 173
digital tokens 74
directories. *See* directory services; domains, directory; folders
Directory Access 136
directory domain administrator 63, 70, 129, 139
directory servers 29
directory services
  See also domains, directory; Open Directory
  Active Directory 129
  benefits of 27–28
  organization of 129
  overview 29, 259
  standalone server 40
disk images
  encrypting data on 120, 123–125
  installing from 36–37
  NetBoot service 38, 255–257
  running Mac OS 9 from 53
disks
  erasing free space 128
  permissions for 45–47
  quotas 138
  startup 111–112
Disk Utility 35, 38, 46–47, 126, 128
`diskutil` tool 35, 38, 128
distributed computing architecture 277–282
DNS (Domain Name System) service 206, 208–213
  profiling 211
  spoofing 210
Dock preferences 97, 147–148
documentation 19–21
Domain Name System. *See* DNS
domains, directory
  See also LDAP; Open Directory
  Active Directory 129
  administrator for 63, 70, 129, 139
  binding of 264–265
  databases 129
  management of 129
  passwords 264
DoS attack (denial of service) 212
duplication of settings 136
DVDs 33, 52–53, 90, 159–161
Dynamic Host Configuration Protocol (DHCP) 206–208, 264

E

EAP-SecurID authentication 202
EFI (Extensible Firmware Interface) 54, 112
email. *See* mail service
encryption
  See also SSL
  AFP 237
  crypt passwords 137, 265, 297
  disk images 123–125
  FileVault 120–123
  server configuration 41
  SSH 185, 191, 236–238
  standards 104
Energy Saver preferences 98–99, 149–150
enterprise applications, JBoss 252–253
erasing data permanently 32, 33–34, 37, 38, 125–128
everyone user category 114
Exposé preferences 93–94
Extensible Firmware Interface (EFI) 54, 112

F

fast user switching 61, 159
files 113–128
  encryption 120–125
  erasing 32, 33–34, 37, 38
  integrity checking tools 292
  OpenSSL 181
  permissions 113–119
  web content 249
file services
  See also AFP; FTP; NFS; share points
  authentication 237–238
  FTP 130–134, 238–240
  NFS 240
file sharing 235–242
file systems, securing 32, 36–37, 39, 125–128
File Transfer Protocol (FTP) 130, 134, 236, 238–240
FileVault 30, 104, 120–123, 161, 194
filters
  junk mail 230, 232
  virus 293
Finder preferences 150–152
fingerprint
  RSA 195
  server 195
firewall service 29, 213–219
firmware password 25, 32, 54–58, 111–112
flags 116, 286
folders 113–119
  flags for 116
  group 138, 139
  home 120–123, 130, 133, 161
  permissions for 113–119
  web content 249
free disk space, erasing 128

FTP (File Transfer Protocol) service 130, 134, 236, 238–240

G

GID (group ID) 138
global file permissions 113–117, 119
global password policy 263
grids, computer 277
group accounts 138–139
  See also groups
group folders 138, 139
groups 114, 117, 138–139

H

hardware, protecting 49–50
help, using 18
HISEC (Highly Secure) templates 129
home folders 120–123, 130, 133, 161
hostconfig entries 284–285
host name, changing 186

I

iChat service 223–225
identities, certificate 173, 174
identity certificates. *See* certificates
images. *See* disk images; NetBoot; Network Install
IMAP (Internet Message Access Protocol) 227
importing certificates 176–177
inherited preferences 142
installation
  auditing tools 284
  command line 37–38
  from removable media 33–34, 52–53
  installer packages 106
  overview 31–32
  permission repair 45–47
  remote 34–37
  server software 33–36, 42–45
  setup 39–42
installer packages 44, 106
`installer` tool 38
install images 38
instant messaging 223–225
internal Software Update server 43
International preferences 99
Internet, client preferences 152–154
Internet Message Access Protocol (IMAP) 227
Internet Printing Protocol (IPP) 267
internet protocol 101, 198, 205
IP addresses
  DHCP 206–208
  DNS recursion 209–210
  firewall service 213–219
  IPv6 notation 205–206
  QTSS 272
IPFilter service 213–219
IP firewall service 213–219
IPP (Internet Printing Protocol) 267
IPv6 addressing 205–206

J

J2EE architecture 252
Jabber instant messaging project 223–225
JBoss application server 252–253
junk mail screening 230, 232

K

KDC (Kerberos Key Distribution Center). *See* Kerberos
Kerberos 130, 201, 264, 296, 302
  authentication 201, 278, 281, 301
  Key Distribution Center (KDC) 28
  mail service 231
  Open Directory 130, 194
  passwords 264, 296
  WebDAV 248
Keyboard preferences 99–100
Keychain Access application 30, 75–79
Keychain certificates 75
keychain services 26, 39, 75–79, 121–122
key item 75
key services 26, 171
known_hosts file 196

L

L2TP (Layer Two Tunneling Protocol) 198–199
LANs (local area networks) 240
layered security architecture 25
Layer Two Tunneling Protocol (L2TP) 198–199
LDAP (Lightweight Directory Access Protocol)
service 263, 295–296
  binding 264–265
  overview 259
  passwords 296
  security policy 264–265
  SSL 262
LDAPv3 access 129
Lightweight Directory Access Protocol. *See* LDAP
Line Printer Remote (LPR) printing 268
local accounts, securing 140
local area networks (LANs) 240
local installation 33–34
local system logging 290
local versus network home folders 130
login 105, 155–159, 296
  keychain 75–76
  passwords 296
  preferences 155–159
  remote 105, 185, 191
  warning banner 61–62
login scripts 155, 156

logs
audit 289
configuration 289–292
firewall service 219
iChat 225
NetBoot 257
SSL 247
LPR (Line Printer Remote) printing 268

M

Mach 24
Mac OS 9 51–53, 90–92
Mailman 227
mail service
authentication 230–231
certificates 228–229
group settings 138
Kerberos 231
key services 171
protocols for 227–230, 234
virus filtering 293
managed accounts 136–146
managed preferences
See also preferences
.Mac 83–84
Accounts 85–87
Applications 145–146
Classic 146–147
Dashboard 93–94
Date & Time 94–95, 187
Desktop 95–97
Dock 97, 147–148
Energy Saver 98–99, 149–150
Finder 150–152
International 99
Internet 152–154
Keyboard 99–100
Login 155–159
Media Access 159–161
Mobility 161–163
Mouse 99–100
Network 100–101, 163–164
Printing 102–103, 165–166
Security 30
Sharing 105
Software Update 167
Sound 107
Spotlight 109–110
Startup Disk 111–112
System 168–169
Universal Access 112, 169–170
managed user accounts 63, 67, 136–146
man-in-the-middle attacks 196
masquerading, IP 220–222
Media Access 159–161
microphone, securing 107
mining, server 211
mobile accounts 161–163, 297
Mobility preferences 161–163
Mouse preferences 99–100
movies, QuickTime cache 103–104
See also streaming media
MS-CHAPv2 authentication 201
multicast 38
multimedia 271–275
multiple users, permissions for 117

N

name server. *See* DNS
naming conventions, computers 40, 186
NAT (Network Address Translation) 220–222
NetBoot service 38, 255–257
NetInfo domains 296
Network Address Translation (NAT) 220–222
Network File System (NFS) 133–134, 236, 240
Network Install 20, 38
Network preferences 163–164
network services
See also IP addresses
DHCP 206–208, 264
directory domains 135
DNS 208–213
FileVault limitations 121
home folders 129
IPv6 addressing 205–206
keychains 79
NTP 41, 94, 184
planning of networks 29–30
preferences 100–101, 163–164
sharing 105
VPN 198–202
wireless preferences 89
Network Time Protocol (NTP) 41, 94, 184
network views 135–136
NFS (Network File System) 133–134, 236, 240
nodes, directory. *See* domains, directory
nonadministrator user accounts 63, 67–69
NSUmask 119
NT Domain services 130, 241, 268
NTP (network time protocol) 41, 94, 184
`nvram` tool 58

O

Open Directory 137, 260, 264, 265, 266, 295–303
See also domains, directory
Active Directory 129
and iChat 223
authentication 295–296
certificates 262–263
configuration 262–264

definition 129
DNS recursion 209
Kerberos 194
overview 27–28, 259
password type 137, 263, 265, 302
replication of 264–265
weblogs 250
Open Directory master 264–266
Open Directory Password Server 260, 264, 296
synchronization of 264
Open Directory replica 264–266
Open Firmware password 32, 54–58, 111
Open Firmware Password application 25, 32, 54–56
OpenLDAP. *See* Open Directory
open source modules
See also Kerberos; Open Directory
Apache 243
Jabber 224
OpenSSL 174, 176, 181
overview 23–25
SpamAssassin 232
optical drives 35–36
overriding preferences 141–142
owners, privileges of 114

P

parental controls 67
partitions, erasing data from 126
Password Assistant 73
passwords 263, 264
administrator 40, 264
authentication setup 72, 73
crypt 137, 297, 301
FileVault master 121
hash type 296
keychain services 26, 75
LDAP 263
Open Directory 296
Open Firmware 25, 32, 54–58, 111
preset 34
protection of server 39
screen saver 50, 95–97, 104
security policy 302
shadow 296
SSL passphrase 247
types 263, 295
Xgrid 279
Password Server. *See* Open Directory Password Server
permissions 113–119
access 24
ACLs 45, 46, 113, 117–118, 131
administrator 29
disk 46–47
everyone 114
files 113–119
group 114
owner 114
root 54
share points 131–133
user 138–139
websites 249–250
piggybacking of service 212
PKI (public key infrastructure) 171
See also certificates
playlists
QuickTime Streaming Server 271
Point-to-Point Tunneling Protocol (PPTP) 200
POP (Post Office Protocol) 227
portable computers
file encryption 121, 123
keychains 79
mobile accounts 161–163
POSIX (Portable Operating System Interface) 46, 113–117
Postfix 227
Post Office Protocol (POP) 227
PPTP (Point-to-Point Tunneling Protocol) 200
`praudit` tool 287–288
preference cache 140
preferences 30, 142
See also managed preferences
accounts 30, 85–87
appearance 88
Bluetooth 89
CDs 90, 159–161
computer 142
DVDs 90
fax 102–103
login 155–159
mail 153
overriding 141–142
QuickTime 103–104
screen saver 50, 95–97
speech recognition 108
time 94–95, 187
user 140
preset passwords 34
presets 136
Print & Fax 102–103
print service 102–103, 165–166, 267
private key 121–122, 172, 173
privileges, administrator 29, 46
See also permissions
profiling, DNS 211
protocols
See also specific protocols
file services 130, 236
mail service 227–232, 234
security settings 188
proxy server settings 163–164

public key certificates. *See* certificates
public key infrastructure (PKI) 171

Q

`qtpaccess` tool 273
`qtpasswd` tool 273
QTSS (QuickTime Streaming Server) 271–275
queues, print 268–269
QuickTime preferences 103–104
QuickTime Streaming Server (QTSS) 271–275
quotas, disk space 138

R

realms. *See* Kerberos; WebDAV; websites, accessing
recent items list 88
recursion, DNS 209–210
relays
  QTSS 274–275
Remote Apple Events 105, 204
remote servers
  Apple Remote Desktop 105
  configuration 41
  installation 34–37
  login 62, 185, 191
  system logging 291
removable media
  FileVault limitations 121
  installation from 33–34, 52–53
  preferences 90, 159–161
replication, Open Directory 260
root permissions 54
RSA fingerprint 195
RSA SecurIDs 74, 201–202

S

SACLs (service access control lists) 188, 225, 238, 240, 296
`scp` tool 192
screening
  See also filters
  junk mail 232
  virus 293
Screen Saver preferences 50, 95–97
Secure Empty Trash option 127
Secure Shell (SSH) 192
  See also SSH
  man-in-the-middle attack 196
Secure Sockets Layer. *See* SSL
SecurID 74, 201–202
security
  overview 23–27
  preferences 104
self-signed certificates 173, 177, 181
serial number, server 34–35
Server Admin application 171, 174, 184, 244, 245, 247
server administrator 63
Server Assistant 36, 39–42
Server Message Block/Common Internet File System. *See* SMB/CIFS
servers
  See also remote servers
  configuration 39–42
  directory 29
  fingerprints 195
  installation 34–36, 42–45
  naming 186
  proxy 163–164
  serial numbers for 34–35
server software 33–45
service access control lists (SACLs) 188, 225, 238, 240, 296
setup procedures. *See* installation
7-pass erase method for data 125–128
SFTP (Secure File Transfer Protocol) 236–238
`sftp` tool 130, 192
shadow passwords 296, 302
shared directory domain 297
shared files 235–242
share points 129, 130–134, 240, 256
Sharing preferences 105
Simple Finder 151
Simple Mail Transfer Protocol (SMTP) 227–228, 230, 234
Simple Network Management Protocol (SNMP) 184
single sign-on (SSO) authentication 278, 303
single-user mode 54
sleep settings, securing 98–99, 149
smart cards 27, 74, 137
SMB/CIFS (Server Message Block/Common Internet File System) protocol
  enabling 241
  IPv6 addressing 206
  printing 268
  security overview 236
  share points 130, 133
SMTP (Simple Mail Transfer Protocol) 227–228, 230, 234
SNMP (Simple Network Management Protocol) 184
Software Update service 42–44, 106, 167
Sound preferences 107
spam 230, 232
Speech preferences 108
spoofing
  ARP 212
  DNS 210
Spotlight 109–110
`srm` tool 127
SSH (Secure Shell Host) 185, 191, 238
`sshd` daemon 192

ssh-keygen tool 194
ssh tool 192
SSL (Secure Sockets Layer) 228–229, 262–263
  certificates 171, 174, 178
  iChat service 224
  mail service 228–229
  Open Directory 262–263
  overview 26
  share points 131
  web service 245–247
standalone server 40
standard user accounts 63, 67
startup, securing 54–58
Startup Disk preferences 111–112
startup disks 111–112
startup image, NetBoot 256
stealth mode 218
streaming media 271–275
sudo tool 52–53, 60, 70, 71–72
synchronization
  mobile account data 161
  passwords 264
  time 41, 184
syslogd configuration file 290
system administrator (root) account 63, 70–71
System Preferences 168–169
  See also managed preferences
system software 42–45, 106

T

target disk mode 111
TCP (Transmission Control Protocol) 211, 213, 221
TCP/IP settings 40
Text to Speech feature 108
35-pass erase method for data 126
time synchronization 41, 184
time zone settings 41, 94–95, 187
TLS (Transport Layer Security) protocol 26
tokens, digital 74
Tomcat 252
Transmission Control Protocol (TCP) 211, 213, 221
Transport Layer Security (TLS) protocol 26
trusted binding 264–265
trust policy services 26
tunneling protocols 198–199, 200
two-factor authentication 74
types of accounts 63–64

U

UCE (unsolicited commercial email) 230, 232
UIDs (user IDs) 64–65
umask settings 119
Universal Access preferences 112, 169–170
UNIX 24
unsolicited mail 230, 232
updating
  Software Update service 42–44, 106, 167
  system software 42–45
user accounts
  See also users
  administrator 39, 69–71, 136, 263
  group 138–139
  in directory domains 136
  Kerberos authentication 303
  managed accounts 63, 67
  mobile 161–163, 297
  passwords 263, 296
  standard 67
user ID 64–65
users 117, 137, 145–146, 188, 296
  See also clients; computer lists; preferences; user accounts; Workgroup Manager
  access control 145–146, 188
  auditing 289
  authentication 259–260, 301–304
  categories 114
  fast user switching 61, 159
  home folders 120–123, 130, 133, 161
  network views 135–136
  passwords 137, 263
  permissions for 117, 138–139
  single-user mode 54
  UIDs 64–65

V

validation, system integrity 283–293
Virtual Private Network (VPN) 198–202
virus screening 293
volumes, securing 32, 36–37, 39, 125–128
VPN (Virtual Private Network) 198–202

W

WebDAV (Web-Based Distributed Authoring and Versioning) 244, 248
weblog service 250–251
web modules 244
WebObjects service 253–254
web service 243–247, 262
websites, accessing 249
Windows services 130, 241, 268
Workgroup Manager 117
  accounts 136–139
  ACL permissions 117
  computer lists 139–140
  directory domains 129
  group account management 138–139
  network views 135–136
  overview 28, 129–130
workgroup preferences 142
  See Workgroup Manager

X

Xgrid 277–282
Xserve servers, installation 35–36

Z

zero-out erase method for data 126
zone transfer, DNS 209