




Mac OS X Server

Server Administration
For Version 10.5 Leopard
Second Edition

 Apple Inc.
© 2008 Apple Inc. All rights reserved.

Under the copyright laws, this manual may not be copied, in whole or in part, without the written consent of Apple.

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.

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

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

Apple, the Apple logo, AirPort, AppleTalk, Final Cut Pro, FireWire, iCal, iChat, iDVD, iMovie, iPhoto, iPod, iTunes, Leopard, Mac, Macintosh, the Mac logo, Mac OS, Panther, PowerBook, Power Mac, QuickTime, SuperDrive, Tiger, Xgrid, Xsan, and Xserve are trademarks of Apple Inc., registered in the U.S. and other countries.

Apple Remote Desktop, Finder, the FireWire logo and Safari are trademarks of Apple Inc.

AppleCare and Apple Store are service marks of Apple Inc., registered in the U.S. and other countries. .Mac is a service mark of Apple Inc.

PowerPC is a trademark of International Business Machines Corporation, used under license therefrom.

UNIX® is a registered trademark of The Open Group.

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.

The product described in this manual incorporates copyright protection technology that is protected by method claims of certain U.S. patents and other intellectual property rights owned by Macrovision Corporation and other rights owners. Use of this copyright protection technology must be authorized by Macrovision Corporation and is intended for home and other limited viewing uses only unless otherwise authorized by Macrovision Corporation. Reverse engineering or disassembly is prohibited.

Apparatus Claims of U.S. Patent Nos. 4,631,603, 4,577,216, 4,819,098 and 4,907,093 licensed for limited viewing uses only.

Simultaneously published in the United States and Canada.

019-1186/2008-02-25

Contents

Preface	11 About This Guide
	11 What's New in Server Admin
	12 What's in This Guide
	12 Using Onscreen Help
	13 Mac OS X Server Administration Guides
	14 Viewing PDF Guides Onscreen
	14 Printing PDF Guides
	15 Getting Documentation Updates
	15 Getting Additional Information
Chapter 1	17 System Overview and Supported Standards
	17 System Requirements for Installing Mac OS X Server
	18 Understanding Server Configurations
	19 Advanced Configuration in Action
	20 Mac OS X Server Leopard Enhancements
	21 Supported Standards
	23 Mac OS X Server's UNIX Heritage
Chapter 2	25 Planning
	25 Planning
	26 Planning for Upgrading or Migrating to Mac OS X Server v10.5
	26 Setting Up a Planning Team
	27 Identifying the Servers You'll Need to Set Up
	27 Determining Services to Host on Each Server
	28 Defining a Migration Strategy
	28 Upgrading and Migrating from an Earlier Version of Mac OS X Server
	29 Migrating from Windows NT
	29 Defining an Integration Strategy
	30 Defining Physical Infrastructure Requirements
	30 Defining Server Setup Infrastructure Requirements
	31 Making Sure Required Server Hardware Is Available
	31 Minimizing the Need to Relocate Servers After Setup
	32 Defining Backup and Restore Policies

32	Understanding Backup and Restore Policies
33	Understanding Backup Types
34	Understanding Backup Scheduling
35	Understanding Restores
36	Other Backup Policy Considerations
37	Command-Line Backup and Restoration Tools

Chapter 3

39	Administration Tools
39	Server Admin
40	Opening and Authenticating in Server Admin
40	Server Admin Interface
41	Customizing the Server Admin Environment
42	Server Assistant
42	Workgroup Manager
43	Workgroup Manager Interface
44	Customizing the Workgroup Manager Environment
44	Directory
45	Directory Interface
46	Directory Utility
46	Server Monitor
48	System Image Management
49	Media Streaming Management
49	Command-Line Tools
50	Xgrid Admin
51	Apple Remote Desktop

Chapter 4

53	Security
53	About Physical Security
54	About Network Security
54	Firewalls and Packet Filters
54	Network DMZ
55	VLANs
55	MAC Filtering
56	Transport Encryption
56	Payload Encryption
57	About File Security
57	File and Folder Permissions
57	About File Encryption
58	Secure Delete
58	About Authentication and Authorization
60	Single Sign-On
60	About Certificates, SSL, and Public Key Infrastructure
61	Public and Private Keys

61	Certificates
62	Certificate Authorities (CAs)
62	Identities
62	Self-Signed Certificates
62	Certificate Manager in Server Admin
64	Readying Certificates
64	Requesting a Certificate From a Certificate Authority
65	Creating a Self-Signed Certificate
65	Creating a Certificate Authority
67	Using a CA to Create a Certificate for Someone Else
68	Importing a Certificate
68	Managing Certificates
69	Editing a Certificate
69	Distributing a CA Public Certificate to Clients
70	Deleting a Certificate
70	Renewing an Expiring Certificate
70	Using Certificates
70	SSH and SSH Keys
71	Key-Based SSH Login
71	Generating a Key Pair for SSH
73	Administration Level Security
73	Setting Administration Level Privileges
73	Service Level Security
74	Setting SACL Permissions
74	Security Best Practices
76	Password Guidelines
76	Creating Complex Passwords
Chapter 5	79 Installation and Deployment
	79 Installation Overview
	81 System Requirements for Installing Mac OS X Server
	81 Hardware-Specific Instructions for Installing Mac OS X Server
	81 Gathering the Information You Need
	82 Preparing an Administrator Computer
	82 About The Server Installation Disc
	83 Setting Up Network Services
	83 Connecting to the Directory During Installation
	83 Installing Server Software on a Networked Computer
	83 About Starting Up for Installation
	84 Before Starting Up
	84 Remotely Accessing the Install DVD
	86 Starting Up from the Install DVD
	86 Starting Up from an Alternate Partition

90	Starting Up from a NetBoot Environment
91	Preparing Disks for Installing Mac OS X Server
98	Identifying Remote Servers When Installing Mac OS X Server
99	Installing Server Software Interactively
99	Installing Locally from the Installation Disc
101	Installing Remotely with Server Assistant
102	Installing Remotely with VNC
103	Using the installer Command-Line Tool to Install Server Software
105	Installing Multiple Servers
106	Upgrading a Computer from Mac OS X to Mac OS X Server
106	How to Keep Current

Chapter 6

107	Initial Server Setup
107	Information You Need
107	Postponing Server Setup Following Installation
108	Connecting to the Network During Initial Server Setup
108	Configuring Servers with Multiple Ethernet Ports
109	About Settings Established During Initial Server Setup
109	Specifying Initial Open Directory Usage
111	Not Changing Directory Usage When Upgrading
111	Setting Up a Server as a Standalone Server
111	Setting Up a Server to Connect to a Directory System
112	Using Interactive Server Setup
113	Setting Up a Local Server Interactively
114	Setting Up a Remote Server Interactively
115	Setting Up Multiple Remote Servers Interactively in a Batch
117	Using Automatic Server Setup
118	Creating and Saving Setup Data
119	Setup Data Saved in a File
120	Setup Data Saved in a Directory
121	Keeping Backup Copies of Saved Setup Data
122	Providing Setup Data Files to Servers
124	How a Server Searches for Saved Setup Data
124	Setting Up Servers Automatically Using Data Saved in a File
127	Setting Up Servers Automatically Using Data Saved in a Directory
129	Determining the Status of Setups
129	Using the Destination Pane for Setup Status Information
130	Handling Setup Failures
130	Handling Setup Warnings
131	Getting Upgrade Installation Status Information
131	Setting Up Services
131	Adding Services to the Server View
131	Setting Up Open Directory

132	Setting Up User Management
132	Setting Up File Services
133	Setting Up Print Service
133	Setting Up Web Service
134	Setting Up Mail Service
134	Setting Up Network Services
135	Setting Up System Image and Software Update Services
135	Setting Up Media Streaming and Broadcasting
135	Setting Up Podcast Producer
136	Setting Up WebObjects Service
136	Setting Up iChat Service
136	Setting Up iCal Service

Chapter 7

137	Management
138	Ports Used for Administration
138	Ports Open By Default
138	Computers You Can Use to Administer a Server
139	Setting Up an Administrator Computer
139	Using a Non-Mac OS X Computer for Administration
140	Using the Administration Tools
140	Opening and Authenticating in Server Admin
141	Adding and Removing Servers in Server Admin
142	Grouping Servers Manually
142	Grouping Servers Using Smart Groups
143	Working with Settings for a Specific Server
145	Changing the IP Address of a Server
146	Changing the Server's Host Name After Setup
146	Changing Server Configuration Type
147	Administering Services
147	Adding and Removing Services in Server Admin
148	Importing and Exporting Service Settings
148	Controlling Access to Services
149	Using SSL for Remote Server Administration
150	Managing Sharing
151	Tiered Administration Permissions
151	Defining Administrative Permissions
152	Workgroup Manager Basics
153	Opening and Authenticating in Workgroup Manager
153	Administering Accounts
153	Working with Users and Groups
154	Defining Managed Preferences
156	Working with Directory Data
156	Customizing the Workgroup Manager Environment

157	Working With Pre-Version 10.5 Computers From Version 10.5 Servers
157	Service Configuration Assistants
157	Critical Configuration and Data Files
161	Improving Service Availability
161	Eliminating Single Points of Failure
162	Using Xserve for High Availability
162	Using Backup Power
163	Setting Up Your Server for Automatic Reboot
164	Ensuring Proper Operational Conditions
164	Providing Open Directory Replication
165	Link Aggregation
166	The Link Aggregation Control Protocol (LACP)
166	Link Aggregation Scenarios
168	Setting Up Link Aggregation in Mac OS X Server
169	Monitoring Link Aggregation Status
170	Load Balancing
171	Daemon Overview
171	Viewing Running Daemons
171	Daemon Control

Chapter 8

173	Monitoring
173	Planning a Monitoring Policy
173	Planning Monitoring Response
174	Server Status Widget
174	Server Monitor
175	RAID Admin
175	Console
175	Disk Monitoring Tools
176	Network Monitoring Tools
177	Notification in Server Admin
178	Monitoring Server Status Overviews Using Server Admin
179	Simple Network Management Protocol (SNMP)
180	Enabling SNMP reporting
180	Configuring snmpd
182	Notification and Event Monitoring Daemons
184	Logging
184	Syslog
185	Directory Service Debug Logging
185	Open Directory Logging
186	AFP Logging
186	Additional Monitoring Aids

Chapter 9	187	Sample Setup
	187	A Single Mac OS X Server in a Small Business
	188	How to Set Up the Server
Appendix	197	Mac OS X Server Advanced Worksheet
Glossary	209	
Index	227	

About This Guide

This guide provides a starting point for administering Mac OS X Leopard Server in advanced configuration mode. It contains information about planning, practices, tools, installation, deployment, and more by using Server Admin.

Server Administration is not the only guide you need when administering advanced mode server, but it gives you a basic overview of planning, installing, and maintaining Mac OS X Server using Server Admin.

What's New in Server Admin

Included with Mac OS X Server v10.5 is Server Admin, Apple's powerful, flexible, full-featured server administration tool. Server Admin is reinforced with improvements in standards support and reliability. Server Admin also delivers a number of enhancements:

- Newly refined and streamlined interface
- Share Point management (functionality moved from Workgroup Manager)
- Event notification
- Tiered administration (delegated administrative permissions)
- Ability to hide and show services as needed
- Easy and detailed server status overviews for one or many servers
- Groups of servers
- Smart Groups of servers
- Ability to save and restore service configurations easily
- Ability to save and restore Server Admin preferences easily

What's in This Guide

This guide includes the following chapters:

- Chapter 1, “System Overview and Supported Standards,” provides a brief overview of Mac OS X Server systems and standards.
- Chapter 2, “Planning,” helps you plan for using Mac OS X Server.
- Chapter 3, “Administration Tools,” is a reference to the tools used to administer servers.
- Chapter 4, “Security,” is a brief guide to security policies and practices.
- Chapter 5, “Installation and Deployment,” is an installation guide for Mac OS X Server.
- Chapter 6, “Initial Server Setup,” provides a guide to setting up your server after installation.
- Chapter 7, “Management,” explains how to work with Mac OS X Server and services.
- Chapter 8, “Monitoring,” shows you how to monitor and log into Mac OS X Server.

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

Using Onscreen Help

You can get task instructions onscreen in Help Viewer while you're managing Leopard Server. You can view help on a server or an administrator computer. (An administrator computer is a Mac OS X computer with Leopard Server administration software installed on it.)

To get help for an advanced configuration of Mac OS X Leopard Server:

- Open Server Admin or Workgroup Manager and then:
 - Use the Help menu to search for a task you want to perform.
 - Choose Help > Server Admin Help or Help > Workgroup Manager Help to browse and search the help topics.

The onscreen help contains instructions taken from *Server Administration* and other advanced administration guides described in “Mac OS X Server Administration Guides,” next.

To see the most recent server help topics:

- Make sure the server or administrator computer is connected to the Internet while you're getting help.

Help Viewer automatically retrieves and caches the most recent server help topics from the Internet. When not connected to the Internet, Help Viewer displays cached help topics.

Mac OS X Server Administration Guides

Getting Started covers installation and setup for standard and workgroup configurations of Mac OS X Server. For advanced configurations, *Server Administration* covers planning, installation, setup, and general server administration. A suite of additional guides, listed below, covers advanced planning, setup, and management of individual services. You can get these guides in PDF format from the Mac OS X Server documentation website:

www.apple.com/server/documentation

This guide...	tells you how to:
<i>Getting Started and Installation & Setup Worksheet</i>	Install Mac OS X Server and set it up for the first time.
<i>Command-Line Administration</i>	Install, set up, and manage Mac OS X Server using UNIX command-line tools and configuration files.
<i>File Services Administration</i>	Share selected server volumes or folders among server clients using the AFP, NFS, FTP, and SMB protocols.
<i>iCal Service Administration</i>	Set up and manage iCal shared calendar service.
<i>iChat Service Administration</i>	Set up and manage iChat instant messaging service.
<i>Mac OS X Security Configuration</i>	Make Mac OS X computers (clients) more secure, as required by enterprise and government customers.
<i>Mac OS X Server Security Configuration</i>	Make Mac OS X Server and the computer it's installed on more secure, as required by enterprise and government customers.
<i>Mail Service Administration</i>	Set up and manage IMAP, POP, and SMTP mail services on the server.
<i>Network Services Administration</i>	Set up, configure, and administer DHCP, DNS, VPN, NTP, IP firewall, NAT, and RADIUS services on the server.
<i>Open Directory Administration</i>	Set up and manage directory and authentication services, and configure clients to access directory services.
<i>Podcast Producer Administration</i>	Set up and manage Podcast Producer service to record, process, and distribute podcasts.
<i>Print Service Administration</i>	Host shared printers and manage their associated queues and print jobs.
<i>QuickTime Streaming and Broadcasting Administration</i>	Capture and encode QuickTime content. Set up and manage QuickTime streaming service to deliver media streams live or on demand.
<i>Server Administration</i>	Perform advanced installation and setup of server software, and manage options that apply to multiple services or to the server as a whole.
<i>System Imaging and Software Update Administration</i>	Use NetBoot, NetInstall, and Software Update to automate the management of operating system and other software used by client computers.
<i>Upgrading and Migrating</i>	Use data and service settings from an earlier version of Mac OS X Server or Windows NT.

This guide...	tells you how to:
<i>User Management</i>	Create and manage user accounts, groups, and computers. Set up managed preferences for Mac OS X clients.
<i>Web Technologies Administration</i>	Set up and manage web technologies, including web, blog, webmail, wiki, MySQL, PHP, Ruby on Rails, and WebDAV.
<i>Xgrid Administration and High Performance Computing</i>	Set up and manage computational clusters of Xserve systems and Mac computers.
<i>Mac OS X Server Glossary</i>	Learn about terms used for server and storage products.

Viewing PDF Guides Onscreen

While reading the PDF version of a guide onscreen:

- Show bookmarks to see the guide's outline, and click a bookmark to jump to the corresponding section.
- Search for a word or phrase to see a list of places where it appears in the document. Click a listed place to see the page where it occurs.
- Click a cross-reference to jump to the referenced section. Click a web link to visit the website in your browser.

Printing PDF Guides

If you want to print a guide, you can take these steps to save paper and ink:

- Save ink or toner by not printing the cover page.
- Save color ink on a color printer by looking in the panes of the Print dialog for an option to print in grays or black and white.
- Reduce the bulk of the printed document and save paper by printing more than one page per sheet of paper. In the Print dialog, change Scale to 115% (155% for *Getting Started*). Then choose Layout from the untitled pop-up menu. If your printer supports two-sided (duplex) printing, select one of the Two-Sided options. Otherwise, choose 2 from the Pages per Sheet pop-up menu, and optionally choose Single Hairline from the Border menu. (If you're using Mac OS X v10.4 or earlier, the Scale setting is in the Page Setup dialog and the Layout settings are in the Print dialog.)

You may want to enlarge the printed pages even if you don't print double sided, because the PDF page size is smaller than standard printer paper. In the Print dialog or Page Setup dialog, try changing Scale to 115% (155% for *Getting Started*, which has CD-size pages).

Getting Documentation Updates

Periodically, Apple posts revised help pages and new editions of guides. Some revised help pages update the latest editions of the guides.

- To view new onscreen help topics for a server application, make sure your server or administrator computer is connected to the Internet and click “Latest help topics” or “Staying current” in the main help page for the application.
- To download the latest guides in PDF format, go to the Mac OS X Server documentation website:
www.apple.com/server/documentation
- An RSS feed listing the latest updates to Mac OS X Server documentation and onscreen help is available. To view the feed use an RSS reader application, such as Safari or Mail:
feed://helposx.apple.com/rss/leopard/serverdocupdates.xml

Getting Additional Information

For more information, consult these resources:

- *Read Me documents*—important updates and special information. Look for them on the server discs.
- *Mac OS X Server website* (www.apple.com/server/macosx)—gateway to extensive product and technology information.
- *Mac OS X Server Support website* (www.apple.com/support/macosxserver)—access to hundreds of articles from Apple’s support organization.
- *Apple Discussions website* (discussions.apple.com)—a way to share questions, knowledge, and advice with other administrators.
- *Apple Mailing Lists website* (www.lists.apple.com)—subscribe to mailing lists so you can communicate with other administrators using email.

System Overview and Supported Standards

1

Mac OS X Server gives you everything you need to provide standards-based workgroup and Internet services — delivering a world-class UNIX-based server solution that's easy to deploy and easy to manage.

This chapter contains information you need to make decisions about where and how you deploy Mac OS X Server. It contains general information about configuration options, standard protocols used, its UNIX roots, and network and firewall configurations necessary for Mac OS X Server administration.

System Requirements for Installing Mac OS X Server

The Macintosh desktop computer or server onto which you install Mac OS X Server v10.5 Leopard must have:

- An Intel or PowerPC G4 or G5 processor, 867 MHz or faster
- Built-in FireWire
- At least 1 gigabyte (GB) of random access memory (RAM)
- At least 10 gigabytes (GB) of available disk space
- A new serial number for Mac OS X Server10.5

The serial number used with any previous version of Mac OS X Server will not allow registration in v10.5.

A built-in DVD drive is convenient but not required.

A display and keyboard are optional. You can install server software on a computer that has no display and keyboard by using an administrator computer. For more information, see “Setting Up an Administrator Computer” on page 139.

Understanding Server Configurations

Mac OS X Server can operate in three configurations: standard, workgroup, and advanced. Servers in advanced configurations are the most flexible and require the most skill to administer. You can customize advanced configurations for a variety of purposes.

An advanced configuration of Mac OS X Server gives the experienced system administrator complete control of service configuration to accommodate a wide variety of business needs. After performing initial setup with Setup Assistant, you use powerful administration applications such as Server Admin and Workgroup Manager, or command-line tools, to configure advanced settings for services the server must provide.

The other two configurations are subsets of the possible services and capabilities of an advanced configuration. They have a simplified administration application, named Server Preferences, and are targeted at more specific roles in an organization.

The workgroup configuration of Mac OS X Server is used for a workgroup in an organization with an existing directory server. A workgroup configuration connects to an existing directory server in your organization and uses the users and groups from the organization's directory in a workgroup server directory.

The standard configuration of Mac OS X Server features automated setup and simplified administration for an independent server in a small organization.

The following table highlights the features and capabilities of each configuration.

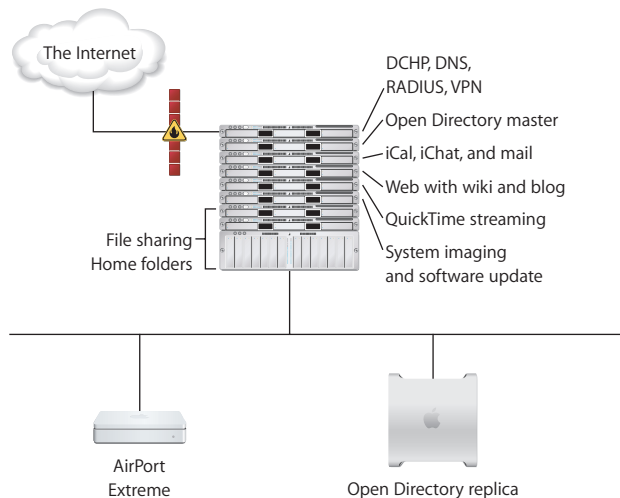
Feature	Advanced	Workgroup	Standard
Service settings changed with...	Server Admin	Server Preferences	Server Preferences
Service settings are...	Unconfigured	Preset to a few common defaults	Preset to common defaults
Users and groups managed with...	Workgroup Manager	Server Preferences	Server Preferences
User service settings automatically provisioned	No	Yes	Yes
Usable as a standalone server	Yes	Yes	Yes
Usable as an Open Directory master	Yes	Yes	Yes
Usable as an Open Directory replica	Yes	Yes	No
Usable as a dedicated network Gateway	Yes	No	Yes

Feature	Advanced	Workgroup	Standard
Usable as an Active Directory plug-in	Yes	Yes	No
Backed up using...	Whatever method implemented by the system administrator	Time Machine preferences pane of System Preferences	Time Machine preferences pane of System Preferences
Dependant on an existing service infrastructure	No	Yes	No
Dependant on an existing well-formed DNS system	Yes	Yes	No

For more information about the Standard and Workgroup configurations and what services are enabled by default for them, see *Getting Started*.

Advanced Configuration in Action

The following illustration depicts several advanced configurations of Mac OS X Server that serve a large organization.



Each server is set up to provide some of the services. For example, one server provides iCal, iChat, and mail service for the organization. Another provides QuickTime media streaming and Podcast Producer.

To ensure high availability of home folders and share points, a master file server and a backup file server have IP failover configured so that if the master fails, the backup transparently takes over. The master and backup file servers use an Xsan storage area network to access the same RAID storage without corrupting it.

For high availability of directory services, Open Directory replicas provide directory service if the Open Directory master goes offline.

The Open Directory domain has user, group, individual computer, and computer group accounts. This allows Mac OS X user preferences to be managed at the group and computer group level.

The web service hosts a website on the Internet for the organization. It also provides wiki websites on the intranet for groups in the organization.

Mac OS X Server Leopard Enhancements

Mac OS X Server includes more than 250 new features, making it the biggest improvement to the server operating system since Mac OS X Server was launched. Here are a few enhancements:

- **Xgrid 2 service:** Xgrid 2 service lets you achieve supercomputer performance levels by distributing computations over collections of dedicated or shared Mac OS X computers. Xgrid 2 features GridAnywhere, allowing Xgrid-enabled software to run where you choose, even if you haven't set up a controller or agents; and Scoreboard for prioritizing which agents are used for each job. Cluster controller provides centralized access to the distributed computing pool, referred to as a computational cluster.
- **File services:** Improved file services includes improved performance and security for each network file service, SMB signing support and secure NFS v3 using Kerberos authentication and AutoFS.
- **iChat Server 2:** iChat Server 2 can federate its community of users with communities of other Extensible Messaging and Presence Protocol (XMPP) messaging systems, such as Google Talk, allowing members of the iChat server community to chat with members of the federated communities.
- **Mail service:** Mail service has added support for mail store clustering when used with Xsan. It also has integrated vacation message functionality. It features improved performance with 64-bit mail services with SMTP, IMAP, and POP.
- **Open Directory 4:** This new version of Open Directory includes new LDAP proxy capability, cross-domain authorization, cascading replication, and replica sets.
- **RADIUS authentication:** RADIUS allows authentication for clients connecting to the network via AirPort Base Stations.
- **QuickTime Streaming Server 6:** Enhanced QuickTime Streaming Server supports 3GPP Release 6 bit-rate adaptation for smooth streaming to mobile phones regardless of network congestion. It integrates with Open Directory on your server when authenticating content delivery, and features improved performance with 64-bit service.

- **Web services:** Web server administrators now have Apache 2.2 (for clean and service upgrade installations) or 1.3 (for upgraded servers). MySQL 5, PHP, and Apache are integrated. Ruby on Rails with Mongrel has been included for simplified development of web-based applications.

Supported Standards

Mac OS X Server provides standards-based workgroup and Internet services. Instead of developing proprietary server technologies, Apple has built on the best open source projects: Samba 3, OpenLDAP, Kerberos, Postfix, Apache, Jabber, SpamAssassin, and more. Mac OS X Server integrates these robust technologies and enhances them with a unified, consistent management interface.

Because it is built on open standards, Mac OS X Server is compatible with existing network and computing infrastructures. It uses native protocols to deliver directory services, file and printer sharing, and secure network access to Mac, Windows, and Linux clients.

A standards-based directory services architecture offers centralized management of network resources using any LDAP server—even proprietary servers such as Microsoft Active Directory. The open source UNIX-based foundation makes it easy to port and deploy existing tools to Mac OS X Server.

The following are some of the standards-based technologies that power Mac OS X Server:

- **Kerberos:** Mac OS X Server integrates an authentication authority based on MIT's Kerberos technology (RFC 1964) to provide users with single sign-on access to secure network resources.

Using strong Kerberos authentication, single sign-on maximizes the security of network resources while providing users with easier access to a broad range of Kerberos-enabled network services.

For services that have not yet been *Kerberized*, the integrated SASL service negotiates the strongest possible authentication protocol.

- **OpenLDAP:** Mac OS X Server includes a robust LDAP directory server and a secure Kerberos password server to provide directory and authentication services to Mac, Windows, and Linux clients. Apple has built the Open Directory server around OpenLDAP, the most widely deployed open source LDAP server, so it can deliver directory services for both Mac-only and mixed-platform environments. LDAP provides a common language for directory access, enabling administrators to consolidate information from different platforms and define one namespace for all network resources. This means a single directory for all Mac, Windows, and Linux systems on the network.

- **RADIUS:** Remote Authentication Dial-In User Service (RADIUS) is an authentication, authorization and accounting protocol used by the 802.1x security standard for controlling network access by clients in mobile or fixed configurations. Mac OS X Server uses RADIUS to integrate with AirPort Base Stations serving as a central MAC address filter database. By configuring RADIUS and Open Directory you can control who has access to your wireless network.

Mac OS X Server uses the FreeRADIUS Server Project. FreeRADIUS supports the requirements of a RADIUS server, shipping with support for LDAP, MySQL, PostgreSQL, Oracle databases, EAP, EAP-MD5, EAP-SIM, EAP-TLS, EAP-TTLS, EAP-PEAP, and Cisco LEAP subtypes. Mac OS X Server supports proxying, with failover and load balancing.

- **Mail Service:** Mac OS X Server uses robust technologies from the open source community to deliver comprehensive, easy-to-use mail server solutions. Full support for Internet mail protocols—Internet Message Access Protocol (IMAP), Post Office Protocol (POP), and Simple Mail Transfer Protocol (SMTP)—ensures compatibility with standards-based mail clients on Mac, Windows, and Linux systems.

- **Web Technologies:** Mac OS X Server web technologies are based on the open source Apache web server, the most widely used HTTP server on the Internet. With performance optimized for Mac OS X Server, Apache provides fast, reliable web hosting and an extensible architecture for delivering dynamic content and sophisticated web services. Because web service in Mac OS X Server is based on Apache, you can add advanced features with plug-in modules.

Mac OS X Server includes everything professional web masters need to deploy sophisticated web services: integrated tools for collaborative publishing, inline scripting, Apache modules, custom CGIs, and JavaServer Pages and Java Servlets. Database-driven sites can be linked to the included MySQL database. ODBC and JDBC connectivity to other database solutions is also supported.

Web service also includes support for Web-based Distributed Authoring and Versioning, known as WebDAV.

- **File Services:** You can configure Mac OS X Server file services to allow clients to access shared files, applications, and other resources over a network. Mac OS X Server supports most major service protocols for maximum compatibility, including:
 - *Apple Filing Protocol (AFP)*, to share resources with clients who use Macintosh computers.
 - *Server Message Block (SMB)*, protocol to share resources with clients who use Windows computers. This protocol is provided by the Samba open source project.
 - *Network File System (NFS)*, to share files and folders with UNIX clients.
 - *File Transfer Protocol (FTP)*, to share files with anyone using FTP client software.

- **IPv6:** IPv6 is short for “Internet Protocol Version 6 (RFC 2460). IPv6 is the Internet’s next-generation protocol designed to replace the current Internet Protocol, IP Version 4 (IPv4, or just IP). IPv6 improves routing and network autoconfiguration. It increases the number of network addresses to over 3×10^{38} , and 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’s network services are fully IPv6 capable and ready to transition to the next generation addressing as well as being fully able to operate with IPv4.
- **SNMP:** Simple Network Management Protocol (SNMP) is used to monitor network-attached devices’ operational status. It is a set of Internet Engineering Task Force (IETF)-designed standards for network management, including an Application Layer protocol, a database schema, and a set of data objects.
Mac OS X Server uses the open source net-snmp suite to provide SNMPv3 (RFCs 3411-3418) service.

Mac OS X Server’s UNIX Heritage

Mac OS X Server has a UNIX-based foundation built around the Mach microkernel and the latest advances from the Berkeley Software Distribution (BSD) open source community. This foundation provides Mac OS X Server with a stable, high-performance, 64-bit computing platform for deploying server-based applications and services.

Mac OS X Server is built on an open source operating system called Darwin, which is part of the BSD family of UNIX-like systems. BSD is a family of UNIX variants descended from Berkeley’s version of UNIX. Also, Mac OS X Server incorporates more than 100 open source projects in addition to proprietary enhancements and extended functionality created by Apple.

The BSD portion of the Mac OS X kernel is derived primarily from FreeBSD, a version of 4.4BSD that offers advanced networking, performance, security, and compatibility features.

In general, BSD variants are derived (sometimes indirectly) from 4.4BSD-Lite Release 2 from the Computer Systems Research Group (CSRG) at the University of California at Berkeley.

Although the BSD portion of Mac OS X is primarily derived from FreeBSD, some changes have been made. To find out more about the low-level changes made, see Apple’s Developer documentation for Darwin.

Before installing and setting up Mac OS X Server do a little planning and become familiar with your options.

The major goals of the planning phase are to make sure that:

- Server user and administrator needs are addressed by the servers you deploy
- Server and service prerequisites that affect installation and initial setup are identified

Installation planning is especially important if you're integrating Mac OS X Server into an existing network, migrating from earlier versions of Mac OS X Server, or preparing to set up multiple servers. But even single-server environments can benefit from a brief assessment of the needs you want a server to address.

Use this chapter to stimulate your thinking. It doesn't present a rigorous planning guide, nor does it provide the details you need to determine whether to implement a particular service and assess its resource requirements. Instead, view this chapter as an opportunity to think about how to maximize the benefits of Mac OS X Server in your environment.

Planning, like design, isn't necessarily a linear process. The sections in this chapter don't require you to follow a mandatory sequence. Different sections in this chapter present suggestions that could be implemented simultaneously or iteratively.

Planning

During the planning stage, determine how you want to use Mac OS X Server and identify whether there's anything you need to accomplish before setting it up.

For example, you might want to convert an existing server to v10.5 and continue hosting directory, file, and mail services for clients on your network.

Before you install server software, you might need to prepare data to migrate to your new server, and perhaps consider whether it's a good time to implement a different directory services solution.

During the planning stage, you'll also decide which installation and server setup options best suit your needs. For example, *Getting Started* contains an example that illustrates server installation and initial setup in a small business scenario with the server in standard configuration mode.

Planning for Upgrading or Migrating to Mac OS X Server v10.5

If you're using a previous version of Mac OS X Server and you want to reuse data and settings, you can upgrade or migrate to v10.5.

You can upgrade to Leopard Server if you're using Mac OS X Server v10.4 Tiger or v10.3 Panther and you don't need to replace server hardware. Upgrading is simple because it preserves existing settings and data. You can perform an upgrade using any of the installation methods described in this chapter or the advanced methods described in this guide.

If you can't perform an upgrade, for example when you need to reformat the startup disk or replace your server hardware, you can migrate data and settings to a computer that you've installed Leopard Server on.

Migration is supported from the latest version of Mac OS X Server v10.4 Tiger, Mac OS X Server v10.3.9 Panther, Mac OS X Server v10.2.8 Jaguar, and Windows NT 4 or later. For complete information about migrating data and settings to a different Mac or Xserve, see *Upgrading and Migrating*. The upgrading and migrating guide provides complete instructions for reusing data and settings in both these scenarios.

Setting Up a Planning Team

Involve individuals in the installation planning process who represent various points of view, and who can help answer the following questions:

- What day-to-day user requirements must a server meet? What activities do server users and workgroups depend on the server for?
If the server is used in a classroom, make sure the instructor who manages its services and administers it daily provides input.
- What user management requirements must be met? Will user computers be diskless and need to be started up using NetBoot? Will Macintosh client management and network home folders be required?
Individuals with server administration experience should work with server users who might not have a technical background, so they'll understand how specific services might benefit them.
- What existing non-Apple services, such as Active Directory, must the server integrate with?

If you've been planning to replace a Windows NT computer, consider using Mac OS X Server with its extensive built-in support for Windows clients. Make sure that administrators familiar with these other systems are part of the planning process.

- What are the characteristics of the network into which the server will be installed? Do you need to upgrade power supplies, switches, or other network components? Is it time to streamline the layout of facilities that house your servers?
An individual with systems and networking knowledge can help with these details as well as completing the "Mac OS X Server Advanced Worksheet" on page 197.

Identifying the Servers You'll Need to Set Up

Conduct a server inventory:

- How many servers do you have?
- How are they used?
- How can you streamline the use of servers you want to keep?
- Are there existing servers that need to be retired? Which ones can Mac OS X Server replace?
- Which non-Apple servers will Mac OS X Server need to be integrated with? Why?
- Do you have Mac OS X Server computers that need to be upgraded to version 10.5?
- How many new Mac OS X Server computers will you need to set up?

Determining Services to Host on Each Server

Identify which services you want to host on each Mac OS X Server and non-Apple server you decide to use.

Distributing services among servers requires an understanding of both users and services. Here are a few examples of how service options and hardware and software requirements can influence what you put on individual servers:

- Directory services implementations can range from using directories and Kerberos authentication hosted by non-Apple servers to setting up Open Directory directories on servers distributed throughout the world. Directory services require thoughtful analysis and planning. *Open Directory Administration* can help you understand the options and opportunities.
- Home folders for network users can be consolidated onto one server or distributed among various servers. Although you can move home folders, you might need to change a large number of user and share point records, so devise a strategy that will persist for a reasonable amount of time. For information about home folders, see *User Management*.

- Some services offer ways to control the amount of disk space used by individual users. For example, you can set up home folder and mail quotas for users. Consider whether using quotas will offer a way to maximize the disk usage on a server that stores home folders and mail databases. *User Management* describes home folder and user mail quotas, and *Mail Service Administration* describes service-wide mail quotas.
- Disk space requirements are also affected by the type of files a server hosts. Creative environments need high-capacity storage to accommodate large media files, but elementary school classrooms have more modest file storage needs. *File Services Administration* describes file sharing.
- If you're setting up a streaming media server, allocate enough disk space to accommodate a certain number of hours of streamed video or audio. For hardware and software requirements and for a setup example, see *QuickTime Streaming and Broadcasting Administration*.
- The number of NetBoot client computers you can connect to a server depends on the server's Ethernet connections, the number of users, the amount of available RAM and disk space, and other factors. DHCP service needs to be available. For NetBoot capacity planning guidelines, see *System Imaging and Software Update Administration*.
- Mac OS X Server offers extensive support for Windows users. You can consolidate Windows user support on servers that provide PDC services, or you can distribute services for Windows users among different servers. The *Open Directory Administration* and *File Services Administration* describe the options available to you.
- If you want to use software RAID to stripe or mirror disks, you'll need two or more drives (they can't be FireWire drives) on a server. For more information, see online Disk Utility Help.

Before finalizing decisions about which servers will host particular services, familiarize yourself with information in the administration guides for services you want to deploy.

Defining a Migration Strategy

If you're using Mac OS X Server v10.2–10.4 or a Windows NT server, examine the opportunities for moving data and settings to Mac OS X Server v10.5.

Upgrading and Migrating from an Earlier Version of Mac OS X Server

If you're using computers with Mac OS X Server versions 10.2, 10.3, or 10.4, consider upgrading or migrating them to an advanced configuration of Mac OS X Server v10.5 Leopard.

If you're using Mac OS X Server v10.4 or v10.3 and you don't need to move to different computer hardware, you can perform an upgrade installation. Upgrading is simple because it preserves your existing settings and data.

When you can't use the upgrade approach, you can migrate data and settings. You'll need to migrate, not upgrade, when:

- A version 10.3 or 10.4 server's hard disk needs reformatting or the server doesn't meet the minimum Leopard Server system requirements. For more information, see "Understanding System Requirements for Installing Mac OS X Server" on page 66.
- You want to move data and settings you've been using on a v10.3 or 10.4 server to different server hardware.
- You want to move data and settings you've been using on a v10.2 server.

Migration is supported from the latest versions of Mac OS X Server v10.4, v10.3, and v10.2. When you migrate, you install and set up an advanced configuration of Leopard Server, restore files onto it from the earlier server, and make manual adjustments as required.

For complete information, read *Upgrading and Migrating*.

Migrating from Windows NT

An advanced configuration of Leopard Server can provide a variety of services to users of Microsoft Windows 95, 98, ME, XP, NT 4, and 2000 computers. By providing these services, Leopard Server can replace Windows NT servers in small workgroups.

For information about migrating users, groups, files, and more from a Windows NT server to Mac OS X Server, see *Upgrading and Migrating*.

Defining an Integration Strategy

Integrating Mac OS X Server into a heterogeneous environment has two aspects:

- Configuring Mac OS X Server to take advantage of existing services
- Configuring non-Apple computers to use Mac OS X Server

The first aspect primarily involves directory services integration. Identify which Mac OS X Server computers will use existing directories (such as Active Directory, LDAPv3, and NIS directories) and existing authentication setups (such as Kerberos). For options and instructions, see *Open Directory Administration*. Integration can be as easy as enabling a Directory Utility option, or it might involve adjusting existing services and Mac OS X Server settings.

The second aspect is largely a matter of determining the support you want Mac OS X Server to provide to Windows computer users. *File Services Administration* and *Open Directory Administration* tell you what's available.

Defining Physical Infrastructure Requirements

Determine whether you need to make site or network topology adjustments before installing and setting up servers.

- Who will administer the server, and what kind of server access will administrators need?

Classroom servers might need to be conveniently accessible for instructors, while servers that host network-wide directory information should be secured with restricted access in a district office building or centralized computer facility.

Because Mac OS X Server administration tools offer complete remote server administration support, there are few times when an administrator should need physical access to a server.

- Are there air conditioning or power requirements that must be met? For this kind of information, see the documentation that comes with server hardware.
- Are you considering upgrading elements such as cables, switches, and power supplies? Now may be a good time to do it.
- Have you configured your TCP/IP network and subnets to support the services and servers you want to deploy?

Defining Server Setup Infrastructure Requirements

The server setup infrastructure consists of the services and servers you set up in advance because other services or servers depend on them.

For example, if you use Mac OS X Server to provide DHCP, network time, or BootP services to other servers, you should set up the servers that provide these services and initiate the services before you set up servers that depend on those services. Or if you want to automate server setup by using setup data stored in a directory, you should set DHCP and directory servers.

The amount of setup infrastructure you require depends on the complexity of your site and what you want to accomplish. In general, DHCP, DNS, and directory services are recommended or required for medium and large server networks:

- The most fundamental infrastructure layer comprises network services like DHCP and DNS.

All services run better if DNS is on the network, and many services require DNS to work properly. If you're not hosting DNS, work with the administrator responsible for the DNS server you'll use when you set up your own servers. DNS requirements for services are published in the service-specific administration guides.

The DHCP setup reflects your physical network topology.

- Another crucial infrastructure component is directory services, required for sharing data among services, servers, and user computers. The most common shared data is for users and groups, but configuration information such as mount records and other directory data is also shared.

A directory services infrastructure is necessary to host cross-platform authentication and when you want services to share the same names and passwords.

Here's an example of the sequence in which you might set up a server infrastructure that includes DNS, DHCP, and directory services. You can set up the services on the same server or on different servers:

- 1 Set up the DNS server.
- 2 Set up DHCP.
- 3 Configure DHCP to specify the DNS server address so it can be served to DHCP clients.
- 4 Set up a directory server, including Windows PDC service if required.
- 5 Populate the directory with data, such as users, groups, and home folder data.

This process can involve importing users and groups, setting up share points, setting up managed preferences, and so forth.

- 6 Configure DHCP to specify the address of the directory server so it can be served to DHCP clients.

Your specific needs can affect this sequence. For example, to use VPN, NAT, or IP Firewall services, you would include their setup into the DNS and DHCP setups.

Making Sure Required Server Hardware Is Available

You might want to postpone setting up a server until all its hardware is in place.

For example, you might not want to set up a server whose data you want to mirror until all the disk drives that you need to set up for mirroring are available. You might also want to wait until a RAID subsystem is set up before setting up a home folder server or other server that will use it.

Minimizing the Need to Relocate Servers After Setup

Before setting up a server, try to place it in its final network location (IP subnet). If you're concerned about preventing unauthorized or premature access during setup, set up a firewall to protect the server while finalizing its configuration.

If you can't avoid moving a server after initial setup, you must change settings that are sensitive to network location before you can use the server. For example, the server's IP address and host name, stored in directories and configuration files on the server, must be updated.

When you move a server, follow these guidelines:

- Minimize the time the server is in its temporary location so the amount of information you need to change is limited.
- Postpone configuring services that depend on network settings until the server is in its final location. Such services include Open Directory replication, Apache settings (such as virtual domains), DHCP, and other network infrastructure settings that other computers depend on.
- Wait to import final user accounts. Limit accounts to test accounts so you minimize the user-specific network information (such as home folder location) that you must change after the move.
- After you move the server, you can change its IP address in the Network pane of System Preferences (or use the `networksetup` tool).

Within a few minutes after you change the server's IP address or name, Mac OS X Server uses the `changeip` command-line tool to update the name, address, other data stored in the Open Directory domain, local directory domain, and service configuration files on the server.

You might need to manually adjust network configurations, such as the server's DNS entries and its DHCP static mapping. For information about the `changeip` tool, see its man page and *Command-Line Administration*.

- Reconfigure the search policy of computers (such as user computers and DHCP servers) that have been configured to use the server in its original location.

Defining Backup and Restore Policies

All storage systems will fail eventually. Either through equipment wear and tear, accident, or disaster, your data and configuration settings are vulnerable to loss. You should have a plan in place to prevent or minimize your data loss.

Understanding Backup and Restore Policies

There are many reasons to have a backup and restore policy. Your data is subject to failure because of failed components, natural or man-made disasters, or data corruption. Sometimes data loss is beyond your control to prevent, but with a backup and restore plan, you can restore your data.

You need to customize backup and restore policies to take into account your situation, your needs, and your own determination of what data needs to be saved, how often, and how much time and effort is used to restore it. Your policy specifies the procedures and practices that fulfill your restoration needs.

Backups are an investment of time, money, and administration effort, and they can affect performance. However, there is a clear return on investment in the form of data integrity. You can avoid substantial financial, legal, and organizational costs with a well-planned, well-executed backup and restore policy.

There are essentially three kinds of restoration needs:

- Restoring a deleted or corrupt file
- Recovering from disk failure (or catastrophic file deletion)
- Archiving data for an organization need (financial, legal, or other need)

Each restoration need determines the type, frequency, and method you use to back up your data.

You might want to keep daily backups of files. This allows for quick restoration of overwritten or deleted files. In such a case you have file-level granularity every day: any single file can be restored the following day.

There are other levels of granularity as well. For example, you might need to restore a full day's data. This is a daily snapshot-level granularity: you can restore your organization's data as it was on a given day.

These daily snapshots might not be practical to maintain every day, so you might choose to keep a set of rolling snapshots that give you daily snapshot-level granularity for only the preceding month.

Other levels of restoration you might want or need could be quarterly or semi-annually.

You might also need archival storage, which is data stored only to be accessed in uncommon circumstances. Archival storage can be permanent, meaning the data is kept for the foreseeable future.

Your organization must determine the following:

- What must be backed up?
- How granular are the restoration needs?
- How often is the data backed up?
- How accessible is the data: in other words, how much time will it take to restore it?
- What processes are in place to recover from a disaster during a backup or restore?

The answers to these questions are an integral part of your backup and restore policy.

Understanding Backup Types

There are many types of backup files (explained below), and within each type are many formats and methods. Each backup type serves a unique purpose and has its own considerations.

- **Full Images:** Full images are byte-level copies of data. They capture the state of the hard disk down to the most basic storage unit. These backups also keep copies of the disk filesystem and the unused or erased portion of the disk in question. They can be used for forensic study of the source disk medium. Such detail often makes file restoration unwieldy. Full Image backups are often compressed and are only decompressed to restore the entire file set.
- **Full File-level Copies:** Full file-level copies are backups that are kept as duplicates. They do not capture the finest detail of unused portions of the source disk, but they do provide a full record of the files as they existed at the time of backup. If a file changes, the next full file-level backup copies the entire data set in addition to the file that changed.
- **Incremental Backups:** Incremental backups start with file-level copies, but they only copy files changed since the last backup. This saves storage space and captures changes as they happen.
- **Snapshots:** Snapshots are copies of data as it was in the past. You can make snapshots from collections of files, or more often from links to other files in a backup file set. Snapshots are useful for making backups of volatile data (data that changes quickly), like databases in use or mail servers sending and receiving mail.

These backup types are not mutually exclusive. They exemplify different approaches to copying data for backup purposes. For example, Mac OS X's Time Machine uses a full file-level copy as a base backup; then it uses incremental backups to create snapshots of a computer's data on a given day.

Understanding Backup Scheduling

Backing up files requires time and resources. Before deciding on a backup plan, consider some of the following questions:

- How much data will be backed up?
- How much time will the backup take?
- When does the backup need to happen?
- What else is the computer doing during that time?
- What sort of resource allocation will be necessary?

For example, how much network bandwidth will be necessary to accommodate the load? How much space on backup drives, or how many backup tapes will be required? What sort of drain on computing resources will occur during backup? What personnel will be necessary for the backup?

You will find that different kinds of backup require different answers to these questions. For example, an incremental file copy might take less time and copy less data than a full file copy (because only a fraction of any given data set will have changed since the last backup).

Therefore an incremental backup might be scheduled during a normal use period because the impact to users and systems may be very low. However, a full image backup might have a very strong impact for users and systems, if done during the normal use period.

Choosing a Backup Rotation Scheme

A backup rotation scheme determines the most efficient way to back up data over a specific period of time. An example of a rotation scheme is the grandfather-father-son rotation scheme. In this scheme, you perform incremental daily backups (son), and full weekly (father) and monthly (grandfather) backups.

In the grandfather-father-son rotation scheme, the number of media sets you use for backup determines how much backup history you have. For example, if you use eight backup sets for daily backups, you have eight days of daily backup history because you'll recycle media sets every eight days.

Understanding Restores

No backup policy or solution is complete without having accompanying plans for data restoration. Depending on what is being restored, you may have different practices and procedures. For example, your organization may have specific tolerances for how long critical systems can be out of use while the data is restored.

You may want to consider the following questions:

- How long will it take to restore data at each level of granularity?
For example, how long will a deleted file or email take to restore? How long will a full hard disk image take to restore? How long would it take to return the whole network to its state three days ago?
- What process is most effective for each type of restore?
For example, why would we roll back the entire server for a single lost file?
- How much administrator action is necessary for each type of restore? How much automation must be developed to best use administrators' time?
- Under what circumstances are the restores initiated? Who and what can start a restore and for what reasons?

Restore practices and procedures must be tested regularly. A backup data set that has not been proven to restore correctly cannot be considered a trustworthy backup. Backup integrity is measured by restore fidelity.

Defining a Backup Verification Mechanism

You should have a strategy for regularly conducting test restorations. Some third-party software providers support this functionality. However, if you're using your own backup solution, you should develop the necessary test procedures.

Other Backup Policy Considerations

Consider the following additional items for your backup policy:

- Should file compression be used? If so, what kind?
- Are there onsite and offsite backups and archives?
- Are there any special considerations for the type of data being stored? For example, for Mac OS X files, can the backup utility preserve file metadata, resource forks, and Access Control List (ACL) privileges?
- Is there sensitive data, such as passwords, social security numbers, phone numbers, medical records, or other legally protected information, that requires special treatment, and must not be backed up without understanding where the data will flow and be stored?

Choosing Backup Media Type

Several factors help you determine what type of media to choose:

- **Cost.** Use cost per GB to determine what media to choose. For example, if your storage needs are limited, you can justify higher cost per GB, but if you need a large amount of storage, cost becomes a big factor in your decision.
One of the most cost-effective storage solutions is a hard drive RAID. Not only does it provide you with a low cost per GB, but it doesn't require the special handling needed by other cost-effective storage types, such as tape drives.
- **Capacity.** If you back up only a small amount of data, low-capacity storage media can do the job. But if you need to back up large amounts of data, use high-capacity devices, such as a RAID.
- **Speed.** When your goal is to keep your server available most of the time, restoration speed becomes a big factor in deciding which type of media to choose. Tape backup systems can be very cost-effective, but they are much slower than a RAID.
- **Reliability.** Successful restoration is the goal of a good backup strategy. If you can't restore lost data, all the effort and cost you spent in backing up data is wasted and the availability of your services compromised.
Therefore, it's important that you choose highly reliable media to prevent data loss. For example, tapes are more reliable than hard disks because they don't contain moving parts.
- **Archive life.** You never know when you'll need your backed up data. Therefore, choose media that is designed to last for a long time. Dust, humidity, and other factors can damage storage media and result in data loss.

Command-Line Backup and Restoration Tools

Mac OS X Server provides several command-line tools for data backup and restoration:

- `rsync`. Use this command to keep a backup copy of your data in sync with the original. The tool `rsync` only copies the files that have changed.
- `ditto`. Use this command to perform full backups.
- `asr`. Use this command to back up and restore an entire volume.

For more information about these commands, see *Command-Line Administration*.

Note: You can use the `launchdctl` command to automate data backup using the aforementioned commands. For more information about using `launchd`, see *Command-Line Administration*.

Manage Mac OS X Server using graphical applications or command-line tools.

Mac OS X Server tools offer diverse approaches to server administration:

- You can administer servers locally (directly on the server you're using) or remotely, from another server, a Mac OS X computer, or a UNIX workstation.
- Graphical applications, such as Server Admin and Workgroup Manager, offer easy-to-use server administration and secure communications for remote server management.

You can use these applications on Mac OS X Server (they're in `/Applications/Server/`) or on a Mac OS X computer where you've installed them, as described in "Setting Up an Administrator Computer" on page 139.

- Command-line tools are available for administrators who prefer to use command-driven server administration.

For remote server management, you can submit commands in a Secure Shell (SSH) session. You can type commands on Mac OS X Server computers and Mac OS X computers using the Terminal application, located in `/Applications/Utilities/`. You can also submit commands from a non-Macintosh computer that's been set up as described in "Using a Non-Mac OS X Computer for Administration" on page 139.

Server Admin

You use Server Admin to administer services on one or more Mac OS X Server computers. Server Admin also lets you specify settings that support multiple services, such as creating and managing SSL certificates, manage file sharing, and specifying which users and groups can access services.

Information about using Server Admin to manage services appears in the individual administration guides and in onscreen information accessible by using the Help menu in Server Admin.

Information about using Server Admin to manage services appears in the individual administration guides and in the following sections.

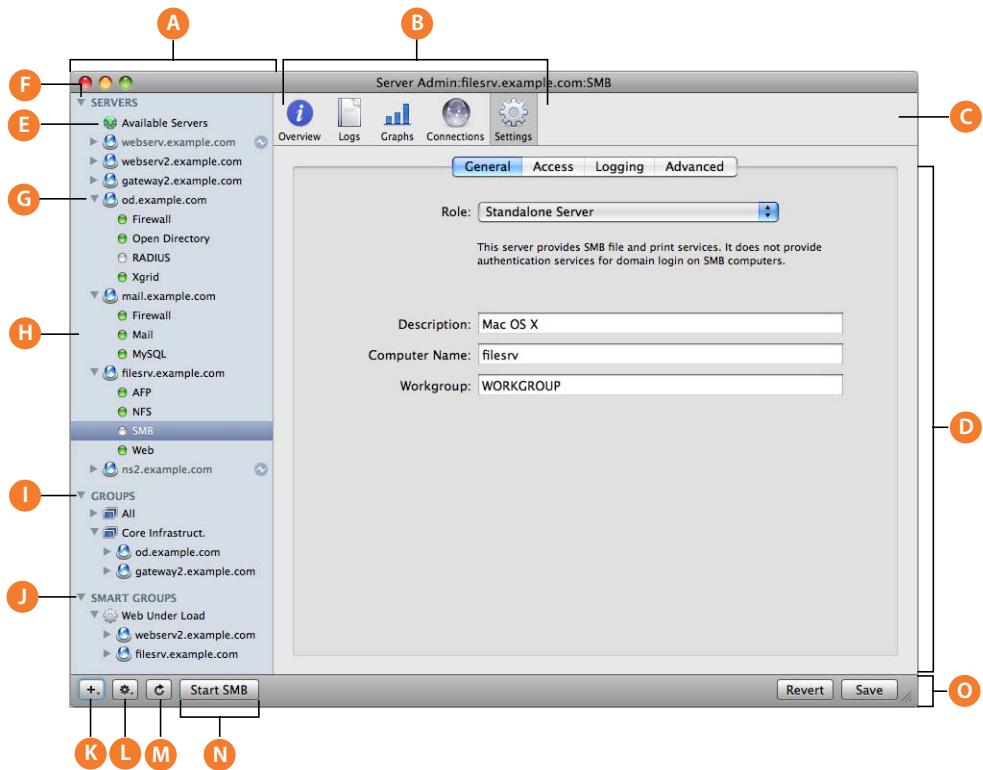
Opening and Authenticating in Server Admin

Server Admin is installed in `/Applications/Server/`, from which you can open it in the Finder. Or you can open Server Admin by clicking the Server Admin icon in the Dock or clicking the Server Admin button on the Workgroup Manager toolbar.

To select a server to work with, enter its IP address or DNS name in the login dialog box, or click Available Servers to choose from a list of servers. Specify the user name and password for a server administrator, then click Connect.

Server Admin Interface

The Server Admin interface is shown here, with each element explained in the following table.



-
- A Server List:** Shows servers, groups, smart groups, and if desired, the administered services for each server
You select a group to view a status summary for all grouped computers.
You select a computer for its overview and server settings.
You select a server's service to control and configure the service.
-
- B Context Buttons:** Shows available information and configuration panes.
-
- C Tool Bar:** Shows available context buttons. If a button is greyed out or can't be clicked, you do not have the administrative permissions to access it.
-
- D Main Work Area:** Shows status and configuration options. This looks different for each service and for each context button selected.
-
- E Available servers:** Lists the local-network scanner, which you can use to discover servers to add to your server list.
-
- F All Servers:** Shows all computers that have been added to Server Admin, regardless of status.
-
- G Server:** Shows the hostname of the managed server. Select to show a hardware, operating system, active service, and system status summary.
-
- H Service:** Shows an administered service for a given server. Select to get service status, logs, and configuration options.
-
- I Group:** Shows an administrator created group of servers. Select to view a status summary for all grouped computers
For more information, see "Grouping Servers Manually" on page 142.
-
- J Smart Group:** Shows an automatic group, populated with servers that meet a predetermined criteria.
For more information, see "Grouping Servers Using Smart Groups" on page 142.
-
- K Add button:** Shows a pop-up menu of items to add to the Server list: servers, groups, and smart groups.
-
- L Action button:** Shows a pop-up menu of actions possible for a selected service, or server, including disconnect server, share the server's screen, and so forth.
-
- M Refresh button:** Allows you to send a status request to all computers visible in the Server list.
-
- N Service Start/Stop button:** When a service is selected, this button allows you to start or stop the service, as appropriate.
-
- O Action bar:** Shows buttons and pop-up menus with commands to act on selected servers or services in the Server list. Click this to save or revert setting changes you've made. this contains the Add button, Action button, service start and stop buttons, and save and revert buttons.
-

Customizing the Server Admin Environment

To control the Server Admin environment, you have the following options.

- To control the list of services to administer, see "Adding and Removing Services in Server Admin" on page 147.
- To control the appearance of Server Admin lists, refresh rates, and other behaviors, choose Server Admin > Preferences.

Server Assistant

Server Assistant is used for:

- Remote server installations
- Initial setup of a local server
- Initial setup of remote servers
- Preparing data for automated setup of an advanced configuration

The Server Assistant initial page is shown here.



Server Assistant is located in `/Applications/Server/`.

For information about using Server Assistant, use its Help buttons, or see Chapter 6, "Initial Server Setup," on page 107.

Workgroup Manager

Mac OS X Server includes Workgroup Manager, a user management tool you can use to create and manage user, group, computer, and computer group accounts. You also use it to access the Inspector, an advanced feature that lets you do raw editing of Open Directory entries.

Workgroup Manager is installed in `/Applications/Server/`, from which you can open it in the Finder. Or you can open Workgroup Manager by clicking `View > Workgroup Manager` in the Server Admin menu bar.

Workgroup Manager works closely with a directory domain. Directory domains are like databases, and are geared towards storing account information and handling authentication.

Information about using Workgroup Manager appears in several documents:

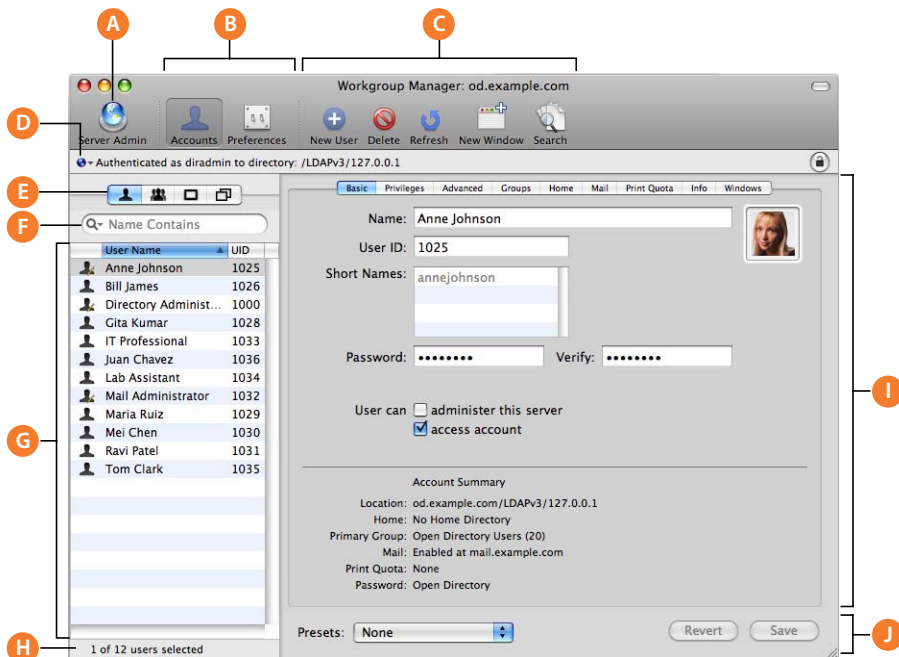
- *User Management* explains how to use Workgroup Manager for account and preference management. This guide also explains how to how to import and export accounts.
- *Open Directory Administration* describes how to use the Inspector.

After opening Workgroup Manager, you can open a Workgroup Manager window by choosing Server > New Workgroup Manager Window.

Important: When connecting to a server or authenticating in Workgroup Manager, make sure the capitalization of the name you enter matches the name of a server administrator or domain administrator account.

Workgroup Manager Interface

The Workgroup Manager interface is shown here, with each element explained in the following table.



A	Server Admin: Click to launch the Server Admin application.
B	Settings Buttons: Click Accounts to view or edit account settings, or click Preferences to view or edit preference settings.
C	Tool Bar: Click the icons to accomplish the various commands. The toolbar is customizable.
D	Directory path: Use to view the directory you are editing. Click the globe icon to select a directory domain. Click the lock to authenticate.
E	Record Type tabs: Use to view records for users, groups, computers, and all records. If the Inspector is enabled, this also contains the Inspector tab.
F	Text filters: Use to enter text to filter record names.
G	Record list display: Use to view all record names for a selected record type.
H	Selection bar: Use to view the number of records found and selected.
I	Main Work Area: Use to work with account, preference, and configuration options. This looks different for each user, group, or preference type.
J	Action zone: Use to save and revert changes, and to make and apply preset configurations to selected records.

Customizing the Workgroup Manager Environment

There are several ways to tailor the Workgroup Manager environment:

- To open Workgroup Manager Preferences, choose Workgroup Manager > Preferences.
You can configure options such as if DNS names are resolved, if the Inspector is enabled, if you need to enter a search query to list records, and what the maximum number of displayed records is.
- To customize the toolbar, choose View > Customize Toolbar.
- To include predefined users and groups in the user and group lists, choose View > Show System Users and Groups.
- To open Server Admin, click the Server Admin toolbar button.

Directory

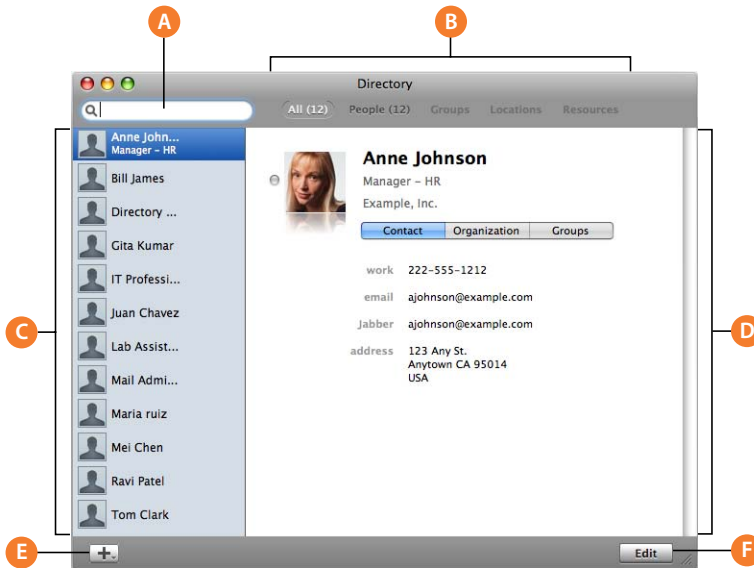
Directory gives users access to shared information about people, groups, locations, and resources within the organization. They can use Directory to share contacts, set up group services, and manage their own contact information.

When users look up information for other people, they'll see more than just contact information. If the person provides a picture, the user will see what he or she looks like. The user can view the person's supervisor and direct reports. The user can see the public groups the person belongs to. The user can also print a map with the person's location pinpointed on it.

Directory takes advantage of several Mac OS X applications. Users can create shared contacts from Address Book entries, click mail addresses to send mail using Mail, or load group web services in Safari.

Directory Interface

The Directory interface is shown here, with each element explained in the following table.

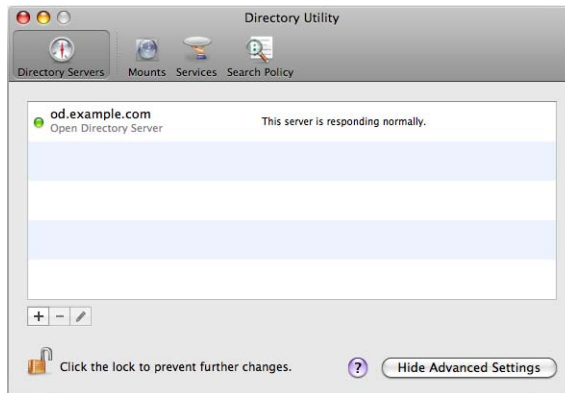


- A Search field:** Use to search record types. Numbers appear to the left of the Record Type buttons to indicate the number of matching records.
- B Record Type buttons:** Click to show the type of directory records desired.
- C Results list:** Use to view the results of the record search.
- D Record view:** Use to view the record selected in the Results list.
- E Add button:** Use to add a person, group, location, or resource record.
- F Edit button:** Click to edit the selected record.

Directory Utility

Directory Utility is the primary application for setting up a Mac OS X computer's connections to Open Directory, Active Directory, and other directory domains, and for defining the computer's search policy and service discovery protocols.

The Directory Utility interface is shown here with advanced configuration options.



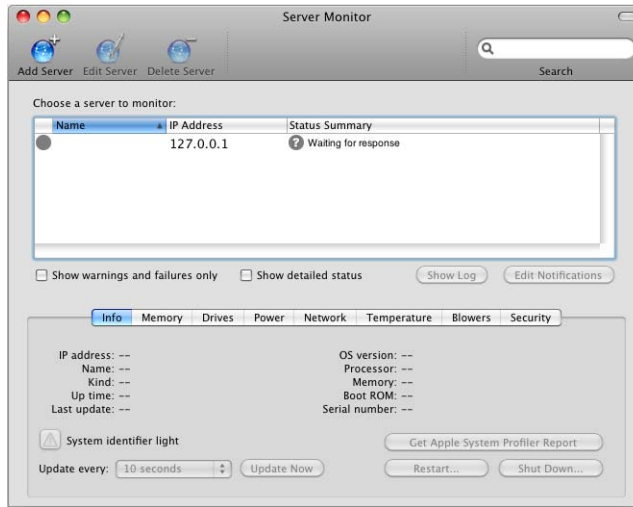
Directory Utility is installed on both Mac OS X Server computers and Mac OS X computers in `/Applications/Utilities/`.

For information about how to use Directory Utility, see *Open Directory Administration* or Directory Utility Help.

Server Monitor

You use Server Monitor to monitor local or remote Xserve hardware and trigger mail notifications when circumstances warrant attention. Server Monitor provides information about the installed operating system, drives, power supply, enclosure and processor temperature, cooling blowers, security, and network.

The Server Monitor interface is shown below.



Server Monitor is installed in `/Applications/Server/` when you install your server or set up an administrator computer. To open Server Monitor, click the Server Monitor icon in the Dock or double-click the Server Monitor icon in `/Applications/Server/`. From within Server Admin, choose `View > Server Monitor`.

To identify the Xserve server to monitor, click `Add Server`, identify the server, and enter user name and password information for an administrator of the server.

To specify how often you want to refresh data, use the “Update every” pop-up menu in the Info pane.

To manage different lists of Xserve servers you want to monitor, choose `File > Export` or `File > Import`. To consolidate lists into one, choose `File > Merge`.

The system identifier lights on the front and back of an Xserve server light when service is required. Use Server Monitor to understand why the lights are on. You can also turn the lights on to identify a particular Xserve server in a rack of servers by selecting the server and clicking “System identifier light” in the Info pane.

To set up Server Monitor to notify you by mail when an Xserve server’s status changes, click `Edit Notifications`. For each server, you set up the conditions for which you want notification. The mail message can come from Server Monitor or from the server.

Server Monitor keeps logs of Server Monitor activity for each Xserve server. To view a log, click `Show Log`. The log shows, for example, Server Monitor attempts to contact the server and whether a connection was successful. The log also shows server status changes. (The logs don’t include system activity on the server.)

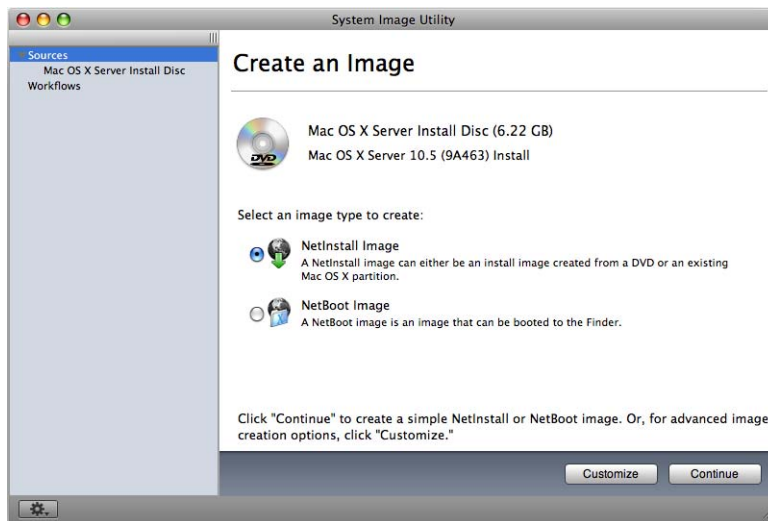
For additional information, see Server Monitor Help.

System Image Management

You can use the following Mac OS X Server applications to set up and manage NetBoot and NetInstall images:

- *System Image Utility* creates Mac OS X disk images. It's installed with Mac OS X Server software in the /Applications/Server/ folder.
- *Server Admin* enables and configures NetBoot service and supporting services. It's installed with Mac OS X Server software in the /Applications/Server/ folder.
- *PackageMaker* creates package files that you use to add software to disk images. Access PackageMaker from Xcode Tools. An installer for Xcode Tools is on the server Install DVD in the Other Installs folder.
- *Property List Editor* edits property lists such as NBImageInfo.plist. Access Property List Editor from Xcode Tools.

The System Image Utility interface is shown below.



System Imaging and Software Update Administration provides instructions for using all these applications.

Media Streaming Management

QuickTime Streaming and Broadcasting Administration provides instructions for administering QuickTime Streaming Server (QTSS) using Server Admin.

QuickTime Streaming and Broadcasting Administration also describes QTSS Publisher, an easy-to-use application for managing media and preparing it for streaming or progressive download.

Command-Line Tools

If you're an administrator who prefers to work in a command-line environment, you can do so with Mac OS X Server.

From the Terminal application in Mac OS X, you can use the built-in UNIX shells (sh, csh, tsh, zsh, bash) to use tools for installing and setting up server software and for configuring and monitoring services. You can also submit commands from a non-Mac OS X computer.

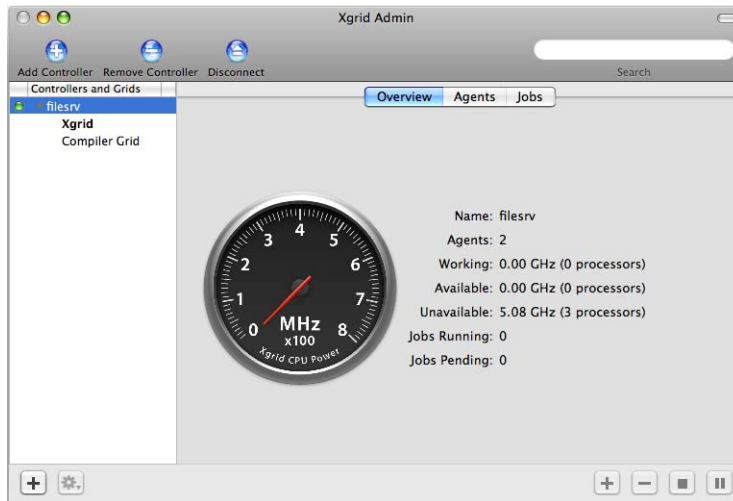
When managing remote servers, you conduct secure administration by working in a Secure Shell (SSH) session.

Command-Line Administration describes Terminal, SSH, server administration commands, and configuration files.

Xgrid Admin

You can use Xgrid Admin to monitor local or remote Xgrid controllers, grids, and jobs. You can add controllers and agents to monitor and specify agents that have not yet joined a grid. You also use Xgrid Admin to pause, stop, or restart jobs.

The System Image Utility interface is shown here.



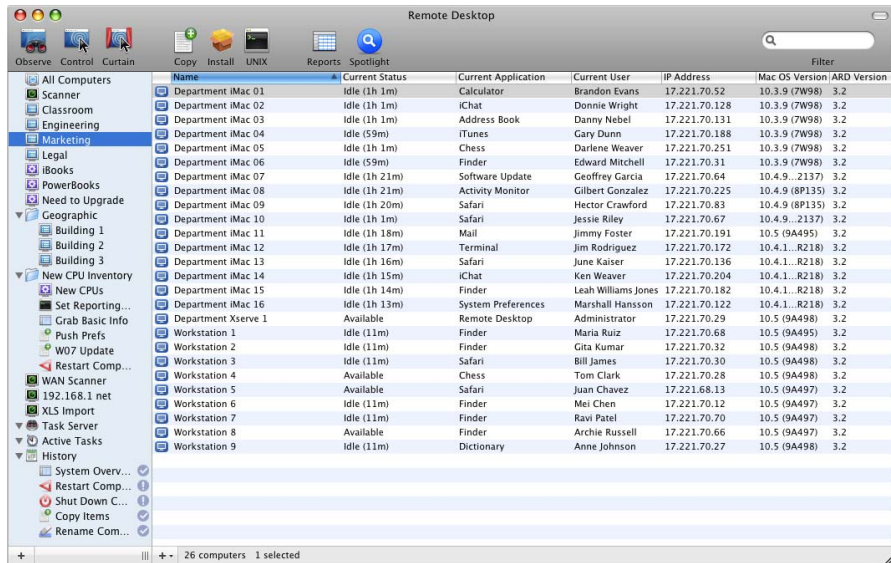
Xgrid Admin is installed in `/Applications/Server/` when you install your server or set up an administrator computer. To open Xgrid Admin, double-click the Xgrid Admin icon in `/Applications/Server/`.

For additional information, see Xgrid Admin help.

Apple Remote Desktop

Apple Remote Desktop (ARD), which you can optionally purchase, is an easy-to-use network-computer management application. It simplifies the setup, monitoring, and maintenance of remote computers and lets you interact with users.

The Apple Remote Desktop interface is shown here.



You can use ARD to control and observe computer screens. You can configure computers and install software. You can conduct one-to-one or one-to-many user interactions to provide help or tutoring. You can perform basic network troubleshooting. And you can generate reports that audit computer hardware characteristics and installed software.

You can also use ARD to control installation on a computer that you start up from an installation disc for Mac OS X Server v10.5 or later, because ARD includes VNC viewer capability.

For more information about Apple Remote Desktop, go to www.apple.com/remotedesktop/.

By vigilantly adhering to security policies and practices, you can minimize the threat to system integrity and data privacy.

Mac OS X Server is built on a robust UNIX foundation that contains many security features in its core architecture. State-of-the-art, standards-based technologies protect your server, network, and data. These technologies include a built-in firewall with stateful packet analysis, strong encryption and authentication services, data security architectures, and support for access control lists (ACLs).

Use this chapter to stimulate your thinking. It doesn't present a rigorous planning outline, nor does it provide the details you need to determine whether to implement a particular security policy and assess its resource requirements. Instead, view this chapter as an opportunity to plan and institute the security policies necessary for your environment.

More information can be found in *Mac OS X Server Security Configuration* and *Mac OS X Security Configuration*.

About Physical Security

The physical security of a server is an often overlooked aspect of computer security. Remember that anyone with physical access to a computer (for example, to open the case, or plug in a keyboard, and so forth) has almost full control over the computer and the data on it. For example, someone with physical access to a computer can:

- Restart the computer from another external disc, bypassing any existing login mechanism.
- Remove hard disks and use forensic data recovery techniques to retrieve data.
- Install hardware-based key-loggers on the local administration keyboard.

In your own organization and environment, you must decide which precautions are necessary, effective, and cost-effective to protect the value of your data and network. For example, in an organization where floor-to-ceiling barriers might be appropriate to protect a server room, securing the air ducts leading to the room might also need to be considered. Other organizations may merely choose a locked server rack or an Open Firmware password.

About Network Security

Network security is as important to data integrity as physical security. Although someone might immediately see the need to lock down an expensive server, he or she might not immediately see the need to restrict access to the data on that same server. The following sections provide considerations, techniques, and technologies to assist you in securing your network.

Firewalls and Packet Filters

Much like a physical firewall that acts as a physical barrier to provide heat and heat damage protection in a building or for a vehicle, a network firewall acts as a barrier for your network assets, preventing data tampering from external sources.

Mac OS X Server's Firewall service is software that protects the network applications running on your Mac OS X Server.

Turning on Firewall service is similar to erecting a wall to limit access. The service scans incoming IP packets and rejects or accepts packets based on the rules you create.

You can restrict access to any IP service running on the server, and you can customize rules for incoming clients or a range of client IP addresses. Services such as Web and FTP services 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 rule. When a packet matches a rule, the action specified in the rule (such as allow or deny) is taken. Then, depending on the action, additional rules might be applied.

Network DMZ

In computer network security, a demilitarized zone (DMZ) is a network area (a subnetwork) that is between an organization's internal network and an external network like the Internet.

You can make connections from the internal and external network to the DMZ, and you can make connections from the DMZ to the external network, but you cannot make connections from the DMZ to the internal network.

This allows an organization to provide services to the external network while protecting the internal network from being compromised by a host in the DMZ. If someone compromises a DMZ host, he or she cannot connect to the internal network.

The DMZ is often used to connect servers that need to be accessible from the external network or Internet, such as mail, web, and DNS servers.

Connections from the external network to the DMZ are often controlled using firewalls and address translation.

You can create a DMZ by configuring your firewall. Each network is connected to a different port on the firewall, called a three-legged firewall setup. This is simple to implement but creates a single point of failure.

Another approach is to use two firewalls with the DMZ in the middle, connected to both firewalls, and with one firewall connected to the internal network and the other to the external network. This is called a screened-subnet firewall.

This setup provides protection in case of firewall misconfiguration, allowing access from the external network to the internal network.

VLANs

Mac OS X Server provides 802.1q Virtual Local Area Network (VLAN) support on the Ethernet ports and secondary PCI gigabit Ethernet cards available or included with Xserves.

VLAN allows multiple computers on different physical LANs to communicate with each other as if they were on the same LAN. Benefits include more efficient network bandwidth utilization and greater security, because broadcast or multicast traffic is only sent to computers on the common network segment. Xserve G5 VLAN support conforms to the IEEE standard 802.1q.

MAC Filtering

MAC Filtering (or layer 2 address filtering) refers to a security access control where a network interface's MAC address, or Ethernet Address (the 48-bit address assigned to each network interface), is used to determine access to the network.

MAC addresses are unique to each card, so using MAC filtering on a network permits and denies network access to specific devices, rather than to specific users or network traffic types. Individual users are not identified by a MAC address, only a device, so an authorized person must have an allowed list of devices that he or she would use to access the network.

In theory, MAC filtering allows a network administrator to permit or deny network access to hosts and devices associated with the MAC address, though in practice there are methods to avoid this form of access control through address modification (spoofing) or the physical exchange of network cards between hosts.

Transport Encryption

Transferring data securely across a network involves encrypting the packet contents sent between two computers. Mac OS X Server can provide Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL) as the cryptographic protocols that provide secure communications on the Internet for such things as web browsing, mail, and other data transfers.

These encryption protocols allow client and server applications to communicate in a way that helps prevent eavesdropping, tampering, and message forgery.

TLS provides endpoint authentication and communications privacy over the Internet using cryptography. These encrypted connections authenticate the server (so its identity is ensured) but the client remains unauthenticated.

To have mutual authentication (where each side of the connection is assured of the identity of the other), use a public key infrastructure (PKI) for the connecting clients.

Mac OS X Server makes use of OpenSSL and has integrated transport encryption into the following tools and services:

- SSH
- VPN
- Web Service
- Mail Service
- Directory Services
- iChat Server
- iCal
- RADIUS

Payload Encryption

Rather than encrypting the transfer of a file across the network, you can encrypt the contents of the file instead. Files with strong encryption might be captured in transit, but would still be unreadable. Most transport encryption requires the participation of both parties in the transaction. Some services (such as SMTP mail service) can't reliably use such techniques, so encrypting the file itself is the only method of reliably securing the file content.

To learn more about file encryption, see "About File Encryption" on page 57.

About File Security

By default, files and folders are owned by the user who creates them. After they're created, items keep their privileges (a combination of ownership and permissions) even when moved, unless the privileges are explicitly changed by their owners or an administrator. Therefore, new files and folders you create are not accessible by client users if they are created in a folder that the users don't have privileges for.

When setting up share points, make sure that items allow appropriate access privileges for the users you want to share them with.

File and Folder Permissions

Mac OS X Server supports two kinds of file and folder permissions:

- Standard Portable Operating System Interface (POSIX) permissions
- Access Control Lists (ACLs)

POSIX permissions let you control access to files and folders based on three categories of users: Owner, Group, and Everyone.

Although these permissions control who can access a file or a folder, they lack the flexibility and granularity that many organizations require to deal with elaborate user environments.

ACL permissions provide an extended set of permissions for files or folders and allow you to set multiple users and groups as owners. In addition, ACLs are compatible with Windows Server 2003 and Windows XP, giving you added flexibility in a multiplatform environment.

For more information about file permissions, see *File Services Administration* and *Mac OS X Server Security Configuration*.

About File Encryption

Mac OS X has a number of technologies that can perform file encryption, including:

- **FileVault:** FileVault performs on-the-fly encryption on each user's home folder. This encrypts the entire directory in one virtual volume, which is mounted and the data is unencrypted as needed.
- **Secure VM:** Secure VM encrypts system virtual memory (memory data temporarily written to the hard disk), not user files. It improves system security by keeping virtual memory files from being read and exploited.

- **Disk Utility:** Disk Utility can create disk images whose contents are encrypted and password protected. Disk images act like removable media such as external hard disks or USB memory sticks, but they exist only as files on the computer. After you create an encrypted disk image, double-click it to mount it. Files you drag onto the mounted image are encrypted and stored on the disk image. You can send this disk image to other Mac OS X users. With the unlocking password, they can retrieve the files you locked in the disk image.

For additional information, the following methods of encrypting files can be found in the *Mac OS X Server Security Configuration Guide*:

- Creating a New Encrypted Disk Image
- Creating an Encrypted Disk Image from Existing Data

Secure Delete

When a file is put in the Trash and the Trash is emptied, or when a file is removed using the `rm` UNIX tool, the files are not removed from disk. Instead, they are removed from the list of files the operating system (OS) tracks and does not write over.

Any space on your hard disk that is free space (places the OS can put a file) most likely contains previously deleted files. Such files can be retrieved using undelete utilities and forensic analysis.

To truly remove the data from disk, you must use a more secure delete method. Security experts advise writing over deleted files and free space multiple times with random data.

Mac OS X Server provides the following tools to allow you to securely delete files:

- Secure Empty Trash (a command in the Finder menu to use instead of “Empty Trash”)
- `srm` (a UNIX utility that securely deletes files, used in place of “`rm`”)

About Authentication and Authorization

Authentication is verifying a person’s identity, but authorization is verifying that an authenticated person has the authority to perform a certain action. Authentication is necessary for authorization.

In a computing context, when you provide a login name and password, you are authenticated to the computer because it assumes only one person (you) knows both the login name and the password. After you are authenticated, the operating system checks lists of people who are permitted to access certain files, and if you are authorized to access them, you are permitted to. Because authorization can’t occur without authentication, authorization is sometimes used to mean the combination of authentication and authorization.

In Mac OS X Server, users trying to use various services (like logging in to a directory-aware workstation, or trying to mount a remote volume) must authenticate by providing a login name and password before any privileges for the users can be determined.

You have several options for authenticating users:

- **Open Directory authentication.** Based on the standard Simple Authentication and Security Layer (SASL) protocol, Open Directory authentication supports many authentication methods, including CRAM-MD5, APOP, WebDAV, SHA-1, LAN Manager, NTLMv2, and Kerberos.

Authentication methods can be selectively disabled to make password storage on the server more secure. For example, if no clients will use Windows services, you can disable the NTLMv1 and LAN Manager authentication methods to prevent storing passwords on the server using these methods. Then someone who somehow gains access to your password database can't exploit weaknesses in these authentication methods to crack passwords.

Open Directory authentication lets you set up password policies for individual users or for all users whose records are stored in a particular directory, with exceptions if required. Open Directory authentication also lets you specify password policies for individual directory replicas.

For example, you can specify a minimum password length or require a user to change the password the next time he or she logs in. You can also disable login for inactive accounts or after a specified number of failed login attempts.

- **Kerberos v5 authentication.** Using Kerberos authentication allows integration into existing Kerberos environments. The Key Distribution Center (KDC) on Mac OS X Server offers full support for password policies you set up on the server. Using Kerberos also provides a feature known as *single sign-on*, described in the next section.

The following services on Mac OS X Server support Kerberos authentication: Apple Filing Protocol (AFP), mail, File Transfer Protocol (FTP), Secure Shell (SSH), login window, LDAPv3, Virtual Private Network (VPN), iChat Server, screen saver, SMB, iCal, and Apache (via the SPNEGO Simple and Protected GSS-API Negotiation Mechanism protocol).

- **Storing passwords in user accounts.** This approach might be useful when migrating user accounts from earlier server versions. However, this approach may not support clients that require certain network-secure authentication protocols, such as APOP.
- **Non-Apple LDAPv3 authentication.** This approach is available for environments that have LDAPv3 servers set up to authenticate users.
- **RADIUS** (an authentication protocol for controlling network access by clients in mobile or fixed configurations). For more information about RADIUS in Mac OS X Server, see *Network Services Administration*.

Single Sign-On

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 use the single sign-on feature, users and services must be Kerberized—configured for Kerberos authentication—and must 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 sign-on.

This server's Kerberized services also use the server's built-in KDC and are automatically configured for single sign-on. 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 was developed at MIT to provide secure authentication and communication over open networks like the Internet. 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 must authenticate only once a day, week, or other period of time, easing authentication loads for users. Mac OS X Server and Mac OS X versions 10.3 through 10.5 support Kerberos version 5.

About Certificates, SSL, and Public Key Infrastructure

Mac OS X Server supports services that use Secure Sockets Layer (SSL) to ensure encrypted data transfer. It uses a Public Key Infrastructure (PKI) system to generate and maintain certificates for use with SSL-enabled services.

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 exchanging secret information in advance.

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 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” on page 61
- “Certificates” on page 61
- “Certificate Authorities (CAs)” on page 62
- “Identities” on page 62

Public and Private Keys

Within a PKI, two digital keys are created: the public key and the private key. The private key isn’t distributed to anyone and is often encrypted by a passphrase. The public key is distributed to other communicating parties.

Basic key capabilities can be summed up as:

Key type	Capabilities
Public	<ul style="list-style-type: none">• Can encrypt messages that can only be decrypted by the holder of the corresponding Private key.• Can verify the signature on a message to ensure that it is coming from a Private key.
Private	<ul style="list-style-type: none">• Can digitally sign a message or certificate, claiming authenticity.• Can decrypt messages that were encrypted with the Public key.• Can encrypt messages that 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 will send its public key to a connecting client and initiate negotiation for a secure connection. The connecting client uses the public key to encrypt a response to the negotiation. The mail server, because 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 issued by a Certificate Authority (CA). A user can digitally sign messages using a private key; then, the receiver can verify the signature using the public key in the CA-issued certificate.

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) that contains:

- The public key half of a public-private key pair
- The key user's identity information, such as a person's 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 a CA, or the key user

Certificate Authorities (CAs)

A CA is an entity that signs and issues digital identity certificates claiming that a party is correctly identified. In this sense, a CA is a trusted third party used by other parties when performing transactions.

In x.509 systems such as Mac OS X, CAs are hierarchical, with CAs being certified by higher CAs, until you reach a root authority. A root authority is a CA that's trusted by the parties, so it doesn't need to be authenticated by another CA. The hierarchy of certificates is top-down, with the root authority's certificate at the top.

A CA can be a company that signs and issues a public key certificate. The certificate attests that the public key belongs to the owner recorded in the certificate.

In a sense, a CA is a digital notary public. You request a certificate by providing the CA with your identity information, contact information, and the public key. The CA then verifies your information so users can trust certificates issued for you by the CA..

Identities

Identities, in the context of the Mac OS X Server Certificate Manager, include signed certificates 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.

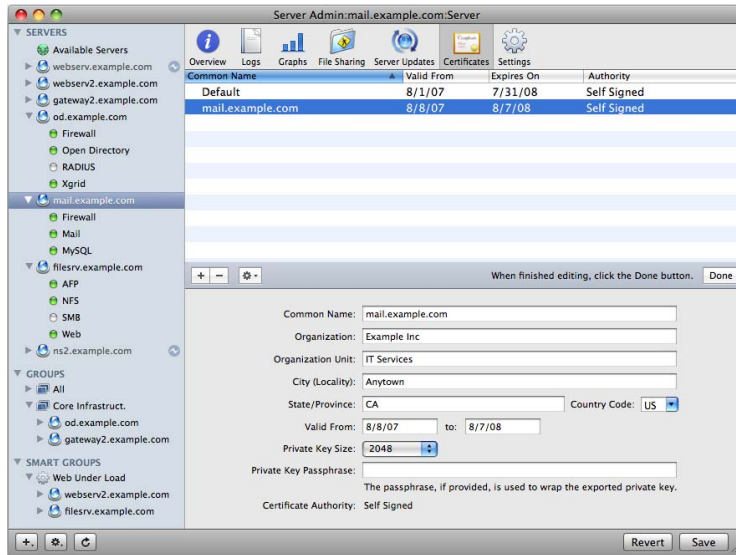
Self-Signed Certificates

Self-signed certificates are certificates that are digitally signed by the private key corresponding to the public key 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.

Certificate Manager in Server Admin

Mac OS X Server's Certificate Manager is integrated into Server Admin to help you create, use, and maintain identities for SSL-enabled services.

The Server Admin interface is shown below, with the Certificate Manager selected.



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 you to create self-signed certificates and certificate-signing requests (CSRs) to obtain certificates signed by a CA. The certificates, self-signed or signed by a CA, are accessible by the services that support SSL.

Identities that were created and stored in OpenSSL files can also be imported into Certificate Manager. They are accessible to services that support SSL.

Certificate Manager in Server Admin doesn't 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 these functions, you can use CA Assistant in Keychain Access (located in /Applications/Utilities/). It provides these capabilities and others for working with x.509 certificates.

Self-signed and CA-issued certificates you created in CA Assistant can be used in Certificate Manager by importing the certificate.

Certificate Manager displays the following for each certificate:

- The domain name that the certificate was issued for
- The dates of validity
- The signing authority (such as the CA entity, or if the certificate is self-signed, it reads "Self-Signed")

Readying Certificates

Before you can use SSL in Mac OS X Server's services, you must create or import certificates. You can create self-signed certificates, generate a Certificate Signing Request (CSR) to send to a CA, or import created certificates.

Select a CA to sign your certificate request. If you don't have a CA to sign your request, consider becoming your own CA, and then import your CA certificates into the root trust database of all your managed machines.

If you're using a self-signed certificate, consider using a self-signed CA to sign a CSR for your service usage, then import the public certificate of your CA into the System keychain on all client computers (if you have control of the computers).

Requesting a Certificate From a Certificate Authority

Certificate Manager helps you create a certificate signing request (CSR) to send to your designated CA.

To request a signed certificate:

1 In Server Admin, select the server that has services that support SSL.

2 Click Certificates.

3 Click the Add (+) button below the Certificates list.

4 Fill out identity information.

The common name is the fully qualified domain name of the server that will use SSL-enabled services.

5 Enter starting and ending validity dates.

6 Select a private key size.

The default is 1024 bits.

7 Enter a passphrase for the private key.

This passphrase should be more secure than a normal password.

It is recommended you use at least 20 characters; include mixed case; include numbers, punctuation, or both, have no characters repeat, and have no dictionary terms.

8 Click the Gear button and choose "Generate Certificate Signing Request"

9 Follow the onscreen directions for requesting a signed certificate from your CA.

For example, you might need to do it online or enter an email address.

10 Click Send Request.

11 Click Done to save the identity information.

When the CA replies to the mail, the CA will include the certificate in the text of the reply.

12 Make sure the Certificate is selected in the Certificates field again.

- 13 Click the Gear button, then choose Add Signed or Renewal Certificate from Certificate Authority.
- 14 Copy the characters from “==Begin CSR==” to “==End CSR==” into the text box.
- 15 Click OK.
- 16 Click Save.

Creating a Self-Signed Certificate

When you create an identity in Certificate Manager, you’re creating a self-signed certificate. Certificate Manager creates a private–public 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 In Server Admin, select the server that has services that support SSL.
- 2 Click Certificates.
- 3 Click the Add (+) button.
- 4 Fill out identity information.

The common name is the fully qualified domain name of the server that will use SSL-enabled services.

- 5 Enter starting and ending validity dates.
- 6 Select a private key size (1024 bits is the default).
- 7 Enter a passphrase for the private key.

This passphrase should be more secure than a normal password.

It is recommended you use at least 20 characters, include mixed case, numbers and punctuation, have no characters repeat, and having no dictionary terms.

- 8 Click Done to save the identity information.
- 9 Click Save.

Creating a Certificate Authority

To sign another user’s certificate, you must create a Certificate Authority (CA). Sometimes a CA certificate is referred to as a root certificate. By signing a certificate with the root certificate, you become the trusted third party in that certificate’s transactions, vouching for the identity of the certificate holder.

If you are a large organization, you might decide to issue or sign certificates for people in your organization to use the security benefits of certificates. However, external organizations may not trust or recognize your signing authority.

To create a CA:

- 1 Start Keychain Access.

Keychain Access is a utility found in the /Applications/Utilities/ directory.

- 2 In the Keychain Access menu, select Certificate Assistant > Create a Certificate Authority.

The Certificate Assistant starts. It will guide you through the process of making the CA.

- 3 Choose to create a Self Signed Root CA.

- 4 Provide the Certificate Assistant with the requested information and click Continue.

You need the following information to create a CA:

- An email address
- The name of the issuing authority (you or your organization)

You also decide if you want to override the defaults, and whether to make this CA the organization's default CA. If you do not have a default CA for the organization, allow the Certificate Assistant to make this CA the default.

In most circumstances, you do not want to override the defaults. If you do not override the defaults, skip to step 16.

- 5 If you choose to override the defaults, provide the following information in the next few screens:

- A unique serial number for the root certificate
- The number of days the CA functions before expiring
- The type of user certificate this CA is signing
- Whether to create a CA website for users to access for CA certificate distribution

- 6 Click Continue.

- 7 Provide the Certificate Assistant with the requested information and click Continue.

You need the following information to create a CA:

- An email address of the responsible party for certificates
- The name of the issuing authority (you or your organization)
- The organization name
- The organization unit name
- The location of the issuing authority

- 8 Select a key size and an encryption algorithm for the CA certificate and then click Continue.

A larger key size is more computationally intensive to use, but much more secure. The algorithm you choose depends more on your organizational needs than a technical consideration. DSA and RSA are strong encryption algorithms. DSA is a United States Federal Government standard for digital signatures. RSA is a more recent advance in algorithms.

9 Select a key size and an encryption algorithm for the certificates to be signed, and then click Continue.

10 Select the Key Usage Extensions you need for the CA certificate and then click Continue.

At a minimum, you must select Signature and Certificate Signing.

11 Select the Key Usage Extensions you need for the certificates to be signed and then click Continue.

Default key use selections are based on the type of key selected earlier in the Assistant.

12 Specify other extensions to add the CA certificate and click Continue.

You must select "Include Basic Constraints" and "Use this certificate as a certificate authority"

13 Specify other extensions to add to the CA certificate and then click Continue.

No other extensions are required.

14 Select the keychain "System" to store the CA certificate.

15 Choose to trust certificates on this computer signed by the created CA.

16 Click Continue and authenticate as an administrator to create the certificate and key pair.

17 Read and follow the instructions on the last page of the Certificate Assistant.

You can now issue certificates to trusted parties and sign certificate signing requests.

Using a CA to Create a Certificate for Someone Else

You can use your CA certificate to issue a certificate to someone else. This is sometimes referred to as signing a Certificate Signing Request (CSR). By doing so you are stating you are a trusted party and can verify the identity of the certificate holder.

Before you can create a certificate for someone, that person must generate a CSR. The user can use the Certificate Assistant to generate the CSR and email the request to you. You then use the CSR's text to make the certificate.

To create a certificate for someone else:

- 1 Start Keychain Access.

Keychain Access is a utility found in the /Applications/Utilities/ directory.

- 2 In the Keychain Access menu, select Certificate Assistant > Create a Certificate for Someone Else as a Certificate Signing Authority.

The Certificate Assistant starts, and guides you through the process of making the CA.

- 3 Drag the CSR and drop it on the target area.
- 4 Choose the CA that is the issuer and sign the request.

Also, you can choose to override the request defaults.

- 5 Click Continue.

If you override the request defaults, provide the Certificate Assistant with the requested information and click Continue.

The Certificate is now signed. The default mail application launches with the signed certificate as an attachment.

Importing a Certificate

You can import a previously generated OpenSSL certificate and private key into Certificate Manager. The items are listed as available in the list of identities and are available to SSL-enabled services.

To import an existing OpenSSL style certificate:

- 1 In Server Admin, select the server that has services that support SSL.
- 2 Click Certificates.
- 3 Click the Import button.
- 4 Enter the existing certificate's file name and path.
Alternately, browse for its location.
- 5 Enter the existing private key file's name and path.
Alternately, browse for its location.
- 6 Enter the private key passphrase.
- 7 Click Import.

Managing Certificates

After you create and sign a certificate, you won't do much more with it. You can use Server Admin to edit certificates before a CA signs them. Except for self-signed certificates, you cannot change certificates after a CA signs them.

If the information a certificate possesses (such as contact information) is no longer accurate, or if you believe the private key is compromised, delete the certificate.

Editing a Certificate

After you add a certificate signature, you can't edit the certificate..

However, you can edit a self-signed certificate. You can modify all fields, including domain name and private key passphrase, private key size, and so forth.

If the identity was exported to disk from the system keychain, you must re-export it.

To edit a certificate:

- 1 In Server Admin, select the server that has services that support SSL.
- 2 Click Certificates.
- 3 Select the Certificate Identity to edit.
It must be a self-signed certificate.
- 4 Click the Edit (/) button.
- 5 Click Edit.

Distributing a CA Public Certificate to Clients

If you're using self-signed certificates, a warning appears in most user applications saying that the certificate authority (CA) is not recognized. Other software, such as the LDAP client, 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 distributed to every client computer that connects 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 nonrewritable media, such as a CD-R. Using nonrewritable media prevents the certificate from being corrupted.
- 2 Open the Keychain Access tool by double-clicking the ca.crt icon where the certificate was copied onto the client computer.
- 3 Add the certificate to the System keychain using Keychain Access.

Alternatively, use the `certtool` command in Terminal:

```
sudo certtool i ca.crt k=/System/Library/Keychains/Systems
```

As a result, any client application (such as Safari or Mail) that verifies certificates using the System keychain recognizes certificates signed by your CA.

Deleting a Certificate

When a certificate has expired or been compromised, you must delete it.

To delete a certificate:

- 1 In Server Admin, select the server that has services that support SSL.
- 2 Click Certificates.
- 3 Select the Certificate Identity to delete.
- 4 Click the Remove (-) button, and select Delete.
- 5 Click Save.

Renewing an Expiring Certificate

All certificates have an expiration date and must be updated periodically.

To renew an expiring certificate:

- 1 Request a new certificate from the CA.
If you are your own CA, create a new one using your own root certificate.
- 2 In Server Admin in the Server list, select the server that has the expiring certificate.
- 3 Click Certificates.
- 4 Select the Certificate Identity to edit.
- 5 Click the action button and select “Add signed or renewed certificate from certificate authority.”
- 6 Paste the renewed certificate into the text field and click OK.
- 7 Click the Edit button to make the certificate editable.
- 8 Adjust the dates for the certificate.
- 9 Click Save.

Using Certificates

In Server Admin, the various services like Web, Mail, VPN, and so on will display a pop-up list of certificates that the administrator can choose from. The services vary in appearance and therefore the pop-up list location varies. Consult the administration guide for the service you’re trying to use with a certificate.

SSH and SSH Keys

SSH is a network protocol that establishes a secure channel between your computer and a remote computer. It uses public-key cryptography to authenticate the remote computer. It also provides traffic encryption and data integrity exchanged between the two computers.

SSH is frequently used to log in to a remote machine to execute commands, but you can also use it to create a secure data tunnel, forwarding through an arbitrary TCP port. You can also use SSH to transfer files using SFTP and SCP. By default, an SSH server uses the standard TCP port 22.

Mac OS X Server uses OpenSSH as the basis for its SSH tools.

Key-Based SSH Login

Key-based authentication is helpful for tasks such as automating file transfers and backups and for creating failover scripts because it allows computers to communicate without a user needing to enter a password. It is not secure to copy the private key of one computer to another computer.

Important: Key-based authentication has risks. If the private key you generate becomes compromised, unauthorized users can access your computers. You must determine whether the advantages of key-based authentication are worth the risks.

Generating a Key Pair for SSH

The following outlines the process of setting up key-based SSH login on Mac OS X and Mac OS X Server. To set up key-based SSH, you must generate the keys the two computers will use to establish and validate the identity of each other. This doesn't authorize all users of the computer to have SSH access. Keys must be generated for each user account.

To do this, run the following commands in Terminal:

- 1 Verify that a `.ssh` folder exists in your home folder by entering the command:

```
ls -ld ~/.ssh.
```

If `.ssh` is listed in the output, move to step 2. If `.ssh` is not listed in the output, run `mkdir ~/.ssh` and continue to step 2.

- 2 Change directories in the shell to the hidden `.ssh` directory by entering the following command:

```
cd ~/.ssh
```

- 3 Generate the public and private keys by entering the following command:

```
ssh-keygen -b 1024 -t rsa -f id_rsa -P ''
```

The `-b` flag sets the length of the keys to 1,024-bits, `-t` indicates to use the RSA hashing algorithm, `-f` sets the file name as `id_rsa`, and `-P` followed by two single-quote marks sets the private key password to be null. The null private key password allows for automated SSH connections. Keys are equivalent to passwords so you should keep them private and protected.

- 4 Copy the public key into the authorized key file by entering the following command:

```
cat id_rsa.pub >> authorized_keys2
```

- 5 Change the permissions of the private key by entering the following command:

```
chmod go-rwx ~/.ssh/.id_rsa
```

The permissions on the private key must be set so the file can only be changed by the group and owner.

- 6 Copy the public key and the authorized key lists to the specified user's home folder on the remote computer by entering the following command:

```
scp authorized_keys2 username@remotemachine:~/.ssh/
```

If you need to establish two-way communication between servers, repeat the above process on the second computer.

This process must be repeated for each user that needs to be able to open a key-based SSH session. The root user is not excluded from this requirement. The home folder for the root user on Mac OS X Server is located at `/var/root/`.

Key-Based SSH with Scripting Sample

A cluster of servers is an ideal environment for using key-based SSH. The following Perl script is a trivial scripting example that should not be implemented. It demonstrates connecting over an SSH tunnel to all servers defined in the variable `serverList`, running `softwareupdate`, installing available updates, and restarting the computer if necessary. The script assumes that key-based SSH has been properly set up for the root user on all servers to be updated.

```
#!/usr/bin/perl

# \@ is the escape sequence for the "@" symbol.
my @serverList = ('root@exampleserver1.example.com',
'root@exampleserver2.example.com');
foreach $server (@serverList) {
open SBUFF, "ssh $server -x -o batchmode=yes 'softwareupdate -i -a' |";

while(<SBUFF>) {
my $flag = 0;
chop($_);
#check for restart text in $_
my $match = "Please restart immediately";
$count = @{{$_ =~ /$match/g}};
if($count > 0) {
$flag = 1;
}
}

close SBUFF;
if($flag == 1) {
`ssh $server -x -o batchmode=yes shutdown -r now`
}
}
```


Administration Level Security

Mac OS X Server can use another level of access control for added security. Administrators can be assigned to services they can configure. These limitations are enacted on a server-by-server basis. This method can be used by an administrator with no restrictions to assign administrative duties to other admin group users. This results in a tiered administration model, where some administrators have more privileges than others for assigned services. This results in a method of access control for individual server features and services.

For example, Alice (the lead administrator) has control over all services on a given server and can limit the ability of other admin group users (like Bob and Cathy) to change settings on the server. She can assign DNS and Firewall service administration to Bob, while leaving mail service administration to Cathy. In this scenario, Cathy can't change the firewall or any service other than mail. Likewise, Bob can't change any services outside of his assigned services.

Tiered administration controls are effective in Server Admin and the `serveradmin` command-line tool. They are not effective against modifying the various UNIX configuration files throughout the system. The UNIX configuration files must be protected with POSIX-type permissions or ACLs.

Setting Administration Level Privileges

You can determine which services other admin group users can modify. To do this, the administrator making the determination must have full, unmodified access.

The process for setting administration level privileges is found in “Tiered Administration Permissions” on page 151.

Service Level Security

You use a Service Access Control List (SACL) to enforce who can use a given service. It is not a means of authentication. It is a list of those who have access rights to use a given service.

SACLs allow you to add a layer of access control on top of standard and ACL permissions.

Only users and groups in a SACL can access its corresponding service. For example, to prevent users from accessing AFP share points on a server, including home folders, remove the users from the AFP service's SACL.

Server Admin in Mac OS X Server allows you to configure SACLs. 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 Apple file service, and so on.

Setting SACL Permissions

SACLs (Service access control lists) allow you to specify which users and groups have access to Mac OS X Server services, including AFP, FTP, and Windows file services.

To set SACL permissions for a service:

- 1 Open Server Admin.
- 2 Select the server from the Servers list.
- 3 Click Settings.
- 4 Click Access.
- 5 To restrict access to all services or deselect this option to set access permissions per service, select “For all services”.
- 6 If you have deselected “For all services,” select a service from the Service list.
- 7 To provide unrestricted access to services, click “Allow all users and groups”.

If you want to restrict access to certain users and groups:

- a Select “Allow only users and groups below.”
 - b Click the Add (+) button to open the Users & Groups drawer.
 - c Drag users and groups from the Users & Groups drawer to the list.
- 8 Click Save.

Security Best Practices

Server administrators must make sure that adequate security measures are implemented to protect a server from attacks. A compromised server risks the resources and data on the server and risks the resources and data on other connected systems. The compromised system can then be used as a base to launch attacks on other systems within or outside your network.

Securing servers requires an assessment of the cost of implementing security with the likelihood of a successful attack and the impact of that attack. It is not possible to eliminate all security risks, but it is possible to minimize risks to efficiently deal with them.

Best Practices for server system administration include, but are not limited to:

- Updating your systems with critical security patches and updates.
- Checking for updates regularly.
- Installing appropriate antivirus tools, using them regularly, and updating virus definition files and software regularly.

Although viruses are far less prevalent on the Mac platform than on Windows, viruses still pose a risk.

- Restricting physical access to the server.
Because local access generally allows an intruder to bypass most system security, secure the server room, server racks, and network junctures. Use security locks. Locking your systems is a prudent thing to do.
- Making sure there is adequate protection against physical damage to servers and ensuring the functioning of the climate control of the server room.
- Taking all additional precautions to secure servers.
For example, enable Open firmware passwords, encrypt passwords where possible, and secure backup media.
- Securing logical access to the server.
For example, remove or disable unnecessary accounts. Accounts for outside parties should be disabled when not in use.
- Configuring SACLs as needed.
Use SACLs to specify who can access services.
- Configuring ACLs as needed.
Use ACLs to control who can access share points and their contents.
- Protecting any account with root or system administrator privileges by following recommended password practices using strong passwords.
For more specific information about passwords, see “Password Guidelines” on page 76 .
- Not using administrator (UNIX “admin” group) accounts for daily use.
Restrict the use of administration privileges by keeping the admin login and password separate from daily use.
- Backing up critical data on the system regularly, with a copy stored at a secure off-site location.
Backup media is of little use in recovery if it is destroyed along with the computer during a machine room fire. Backup/Recovery contingency plans should be tested to ensure that recovery actually works.
- Reviewing system audit logs regularly and investigating unusual traffic.

- Disabling services that are not required on your system.
A vulnerability that occurs in any service on your system can compromise the entire system. In some cases, the default configuration (out of the box) of a system leads to exploitable vulnerabilities in services that were enabled implicitly.
Turning on a service opens up a port from which users can access your system. Although enabling Firewall service helps fend off unauthorized access, an inactive service port remains a vulnerability that an attacker might be able to exploit.
- Enabling Firewall service on servers, especially at the network frontier.
Your server's firewall is the first line of defense against unauthorized access. For more information, see the chapter on setting up Firewall service in *Network Services Administration*. Consider also a third-party hardware firewall as an additional line of defense if your server is highly prone to attack.
- If needed, installing a local firewall on critical or sensitive servers.
Implementing a local firewall protects the system from an attack that might originate from within the organization's network or from the Internet.
- For additional protection, implementing a local Virtual Private Network (VPN) that provides a secure encrypted tunnel for all communication between a client computer and your server application. Some network devices provide a combination of functions: firewall, intrusion detection, and VPN.
- Administering servers remotely.
Manage your servers remotely using applications like Server Admin, Server Monitor, RAID Admin, and Apple Remote Desktop. Minimizing physical access to the systems reduces the possibility of mischief.

Password Guidelines

Many applications and services require that you create passwords to authenticate. Mac OS X includes applications that help create complex passwords (using Password Assistant), and securely store your passwords (using Keychain Access).

Creating Complex Passwords

Use the following tips to create complex passwords:

- Use a mix 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.
- Don't append a number to an alphabetic word (for example, "wacky2") to fulfill the constraint of having a number.
- Don't substitute "look alike" numbers or symbols for letters (for example, "GR33N" instead of "GREEN").
- Don't use proper names.

- Don't use dates.
- Create a password of at least 12 characters. Longer passwords are generally more secure than shorter passwords.
- Use passwords that can't be guessed even by someone who knows you and your interests well.
- Create as random a password as possible.

You can use Password Assistant (located in /System/Library/CoreServices/ to verify the complexity of your password.

Whether you install Mac OS X Server on a single server or a cluster of servers, there are tools and processes to help the installation and deployment succeed.

Some computers come with Mac OS X Server software already installed. Other computers need to have the server software installed. For example, installing Leopard Server on a computer with Mac OS X makes the computer a server with Mac OS X Server.

Installing Leopard Server on an existing server with an Mac OS X Server v10.2–10.4 upgrades the server software to v10.5. If Leopard Server is already installed, installing it again refreshes the server environment.

This chapter includes instructions for a fresh installation of Leopard Server using a variety of methods.

Installation Overview

You've already planned and decided how many and what kind of servers you are going to install.

Step 1: Confirm you meet the requirements

Make sure your target server meets the minimum system requirements. For more information see:

- "System Requirements for Installing Mac OS X Server" on page 81
- "Hardware-Specific Instructions for Installing Mac OS X Server" on page 81

Step 2: Gather your information

Gather all the information you need before you begin. This not only helps to make sure the installation goes smoothly, but it can help you make certain planning decisions. For further information, see:

- Chapter 2, "Planning," on page 25
- Appendix , "Mac OS X Server Advanced Worksheet," on page 197

- “About The Server Installation Disc” on page 82

Step 3: Set up the environment

If you are not in complete control of the network environment (DNS servers, DHCP server, firewall, and so forth) you need to coordinate with your network administrator before installing. A functioning DNS system, with full reverse lookups, and a firewall to allow configuration constitute a bare minimum for the setup environment. If you are planning on connecting the server to an existing directory system, you also need to coordinate efforts with the directory administrator. See the following:

- “Connecting to the Directory During Installation” on page 83
- “Installing Server Software on a Networked Computer” on page 83

If you are administering the server from another computer, you must create an administration computer. For more information, see “Preparing an Administrator Computer” on page 82.

Step 4: Start up the computer from an installation disk

You can't install onto the disk the computer is booted from, but you can upgrade. For clean installations and upgrades, you must start up the server from an installation disk, not from the target disk. See the following:

- “About Starting Up for Installation” on page 83
- “Remotely Accessing the Install DVD” on page 84
- “Starting Up from the Install DVD” on page 86
- “Starting Up from an Alternate Partition” on page 86
- “Starting Up from a NetBoot Environment” on page 90

Step 5: Prepare the target disk

If you are doing a clean installation, you must prepare the target disk by making sure it has the right format and partition scheme. See the following:

- “Preparing Disks for Installing Mac OS X Server” on page 91
- “Choosing a File System” on page 91
- “Partitioning a Hard Disk” on page 93
- “Creating a RAID Set” on page 94
- “Erasing a Disk or Partition” on page 97

Step 6: Start the installer

The installer application takes software from the startup disk and server software packages and installs them on the target disk. See the following:

- “Identifying Remote Servers When Installing Mac OS X Server” on page 98
- “Installing Server Software Interactively” on page 99
- “Installing Locally from the Installation Disc” on page 99
- “Installing Remotely with Server Assistant” on page 101

- “Installing Remotely with VNC” on page 102
- “Using the installer Command-Line Tool to Install Server Software” on page 103

Step 7: Set up services

Restart from the target disk to proceed to setup. For more information about server setup, see “Initial Server Setup” on page 107.

System Requirements for Installing Mac OS X Server

The Macintosh desktop computer or server where you install Mac OS X Server v10.5 Leopard must have:

- An Intel or PowerPC G4 or G5 processor, 867 MHz or faster
- Built-in FireWire
- At least 1 gigabyte (GB) of random access memory (RAM)
- At least 10 gigabytes (GB) of disk space available
- A new serial number for Mac OS X Server 10.5.

The serial number used with any previous version of Mac OS X Server will not allow registration in v10.5.

A built-in DVD drive is convenient but not required.

A display and keyboard are optional. You can install server software on a computer that has no display and keyboard by using an administrator computer. For more information, see “Preparing an Administrator Computer” on page 82.

If you’re using an installation disc for Mac OS X Server v10.5 or later, you can control installation from another computer using VNC viewer software. Open source VNC viewer software is available. Apple Remote Desktop, described on page 51, includes VNC viewer capability.

Hardware-Specific Instructions for Installing Mac OS X Server

When you install server software on Xserve systems, the procedure you use when starting the computer for installation is specific to the kind of Xserve hardware you have. You may need to refer to the *Xserve User’s Guide* or *Xserve Setup Guide* that came with your Xserve, where these procedures are documented.

Gathering the Information You Need

Use the “Mac OS X Server Advanced Worksheet” to record information for each server you want to install. The information below provides supplemental explanations for items on the “Mac OS X Server Advanced Worksheet”. The “Mac OS X Server Advanced Worksheet” is located in the appendix on page 197.

Preparing an Administrator Computer

You can use an administrator computer to install, set up, and administer Mac OS X Server on another computer. An administrator computer is a computer with Mac OS X v10.5 Leopard or Mac OS X Server Leopard that you use to manage remote servers.

When you install and set up Mac OS X Server on a computer that has a display and keyboard, it's already an administrator computer. To make a computer with Mac OS X into an administrator computer, you must install additional software.

Important: If you have administrative applications and tools from Mac OS X Server v10.4 Tiger or earlier, do not use them with Leopard Server.

To enable remote administration of Mac OS X Server from a Mac OS X computer:

- 1 Make sure the Mac OS X computer has Mac OS X v10.5 Leopard installed.
- 2 Make sure the computer has at least 1 GB of RAM and 1 GB of unused disk space.
- 3 Insert the Administration Tools CD.
- 4 Open the Installers folder.
- 5 Open ServerAdministrationSoftware.mpkg to start the Installer, and then follow the onscreen instructions.

About The Server Installation Disc

You can install the server software using the Mac OS X Server Install Disc. This installation disc contains everything you must install Mac OS X Server. It also contains an Other Installs folder, which has installers for upgrading a Mac OS X computer to Mac OS X Server and for separately installing server administration software, the Directory application, the Podcast Capture application, X11 software, and Xcode developer tools.

In addition to the installation disc, Mac OS X Server includes the Administration Tools CD. You use this disc to set up an administrator computer. This disc also contains installers for the Directory application, the Podcast Capture application, and the QTSS Publisher application. For advanced administrators, this disc contains installers for PackageMaker and Property List Editor.

Setting Up Network Services

Before you can install, you must set up or have the following settings for your network service:

- **DNS:** You must have a fully qualified domain name for each server's IP address in the DNS system. The DNS zone must have the reverse-lookup record for the name and address pair. Not having a stable, functioning DNS system with reverse lookup leads to service failures and unexpected behaviors.
- **DHCP:** It is not recommended to assign dynamic IP addresses to servers. If your server gets its IP address through DHCP, set up a static mapping in the DHCP server, so your server gets (via its Ethernet address) the same IP address every time.
- **Firewall or routing:** In addition to any firewall running on your server, the subnet router may have certain network traffic restrictions in place. Make sure your server's IP address is available for the traffic you are planning to handle and the services you are planning to run.

Connecting to the Directory During Installation

If you want to use a server as an Open Directory master, make sure it has an active Ethernet connection to a secure network before installation and initial setup.

Installing Server Software on a Networked Computer

When you start up a computer from a server installation disc, SSH starts so that remote installations can be performed.

Important: Before you install or reinstall Mac OS X Server, make sure the network is secure because SSH gives others access to the computer over the network. For example, design the network topology so you can make the server computer's subnet accessible only to trusted users.

About Starting Up for Installation

The computer can't install to its own startup volume, so you must start up in some other way, such as:

- Optical Media, DVDs
- Alternate volumes (second partitions on the hard disk, or external FireWire disks)
- Netboot

The computer must install from the same disk or image that started up the computer. Mounting another share point with an installer won't work. The installer uses some of the files currently active in the booted system partition for the new installation.

Before Starting Up

If you're performing a clean installation rather than upgrading an existing server, back up any user data that's on the disk or partition where you'll install the server software.

If you're upgrading an existing server, make sure that saved setup data won't be inadvertently detected and used to automatically set up an advanced configuration. Server Assistant looks for saved setup data on all mounted disks and in all directories the server is configured to access. The saved setup data will overwrite the server's existing settings.

For more information about automatic server setup, see "Using Automatic Server Setup" on page 117.

Remotely Accessing the Install DVD

When used as the startup disc, the Install DVD provides some services for remote access. After you start up from the DVD, both SSH and VNC are available for use. VNC enables you to use a VNC viewer (like Apple Remote Desktop) to view the user interface as if you were using the remote computer's keyboard, mouse, and monitor. All the things you could do at the computer using the keyboard and mouse are available remotely, as well as locally. This excludes hard resets, other hardware manipulation, or holding down keys during startup.

SSH enables you to have command-line access to the computer with administrator privileges.

To access the computer with VNC:

- 1 Start the target computer from the Install DVD for Mac OS X Server v10.5 or later. The procedure you use depends on the target server hardware.

To learn more about startup disk options, see "About Starting Up for Installation" on page 83.

- 2 Use your VNC viewer software to open a connection to the target server.
- 3 Identify the target server.

If the VNC viewer includes the target server in a list of available servers, select it in the list. Otherwise, enter an IP address in IPv4 format (000.000.000.000).

If you don't know the IP address and the remote server is on the local subnet, you can use the `sa_srchr` command to identify computers on the local subnet where you can install server software. Enter the following from an existing computer with Mac OS X Server Tools installed:

```
/System/Library/Serverssetup/sa_srchr 224.0.0.1
```

This command returns the IP address and the EthernetID (in addition to other information) of servers on the local subnet that started up from the installation disk.

- 4 When prompted for a password, enter the first eight digits of the server's built-in hardware serial number.

To find a server's serial number, 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.

If you're using Apple Remote Desktop as a VNC viewer, enter the password but don't specify a user name.

To access the computer with SSH:

- 1 Start the target computer from the Install DVD for Mac OS X Server v10.5 or later.

The procedure you use depends on the target server hardware.

To learn more about startup disk options, see "About Starting Up for Installation" on page 83.

- 2 Use the Terminal to open a secure shell connection to the target server.

The user name is root and the password is the first eight digits of the server's built-in hardware serial number.

To find a server's serial number, 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.

If you don't know the IP address and the remote server is on the local subnet, you can use the `sa_srchr` command to identify computers on the local subnet where you can install server software. Enter the following from an existing computer with Mac OS X Server Tools installed:

```
/System/Library/Serverssetup/sa_srchr 224.0.0.1
```

This command will return the IP address, and the EthernetID (in addition to other information) of servers on the local subnet which have started up from the installation disk.

Starting Up from the Install DVD

This is the simplest method of starting the computer, if you have physical access the server, and it has an optical drive.



If the target server is an Xserve with a built-in DVD drive, start the server using the Install DVD by following the instructions in the *Xserve User's Guide* for starting from a system disc.

If the target server has no built-in DVD drive, you can use an external FireWire DVD drive. You can also install server software on an Xserve system that lacks a DVD drive by moving its drive module to another Xserve system that has a DVD drive.

To start up the computer with the installation disc.

- 1 Turn on the computer and insert the Mac OS X Server installation disc into the DVD drive.
- 2 If you're using a built-in DVD drive, restart the computer while holding down the C key. You can release the C key when you see the Apple logo.
Alternatively, you can restart the computer by holding down the Option key, selecting the icon representing the installation disc, and then clicking the right arrow.
You must use this method if you are starting up from an external DVD drive.
- 3 If you're installing on an Xserve, the procedure for starting up from a DVD may be different. For more information, see the *Xserve User's Guide* or *Xserve Setup Guide* that came with your Xserve.
- 4 After the computer restarts, choose the language you want to use during installation and then click the arrow button.

The Installer is now running.

Starting Up from an Alternate Partition

For a single server installation, preparing to start up from an alternate partition can be more time-consuming than simply using the Install DVD. The time required to image, scan, and restore the image to a startup partition may exceed the time taken to install once from the DVD. However, if you are reinstalling regularly, or if you are creating an external Firewire drive-based installation to take to various computers, or if you need some other kind mass distribution (such as clustered Xserves without DVD drives installed), this method can be very efficient.

This method is well suited to installing on computers that you may not have easy physical access to. With sufficient preparation, this method can be modified for easy mass deployment of appropriately licensed copies of Mac OS X Server.

To use this method, you must have an existing installation of some kind on the computer in order to use this method. It is intended for environments where a certain level of existing infrastructure of Mac OS X Server is present, and may be unsuitable for a first server installation. To start from an alternate partition, there are four basic steps.

Step 1: Prepare the disks and partitions on the target computer

Before you proceed, you must have at least two partitions on the target computer. The first is going to be the initial and the final startup partition; the second is the temporary installer partition. You can use a single disk with multiple partitions, or you can use multiple disks. You use Disk Utility to prepare the disks.

For more information about preparing and partitioning a hard disk, see the Disk Utility help.


Step 2: Create a restorable image of the Install DVD

This step doesn't need to be done on the target computer. It can be done on an administrator computer, but there must be enough free space to image the entire Install DVD.

To create an image of the Install DVD:

- 1 Insert the Install DVD.
- 2 Launch Disk Utility.
- 3 Select the first session icon under the optical drive icon.
This is in the list of devices on the left side of the window.
- 4 Select File > New > Disk Image from <device>.
- 5 Give the image a name, select Read-only, Read/Write, or Compressed as the image type, and then click Save.
- 6 After the image is complete, select the image from list on the left.
- 7 In the menu, select Images > Scan Images for Restore.
- 8 Provide an administrator login and password as needed.

The installer disk image can now be restored to your extra partition.

-  **Tip:** If you prefer to use the command-line, you can use `hdiutil` to create the disk image, and `asr` to scan the image for restore. All commands must be done with super-user or root privileges.

For example, this command creates a disk image “Installer.dmg” from the device at disk1s1:

```
hdiutil create -srcdevice disk1s1 Installer.dmg
```

This command scans the image “Installer.dmg” and readies it for restore:

```
asr imagescan --source Installer.dmg
```

Step 3: Restore the image to the alternate partition

You can restore the disk image to a partition within the computer or to an external hard disk. When complete, the newly restored partition functions like the Install DVD. Make sure the alternate partition is at least the size of the disk image.

Restoring the disk image to the partition will erase all existing data on the partition.

To restore the image:

- 1 Start up the target computer.
- 2 Make sure the image does not reside on the partition that is to be erased.
- 3 Launch Disk Utility.
- 4 In the list of devices on the left side of the window, select the installer DVD image.
- 5 Click Restore.
- 6 Drag the installer image from the left side of the window to the Source field.
- 7 Drag the alternate partition from the list of devices on the left side of the window to the Destination field.
- 8 Select Erase Destination.
- 9 Click Restore.

If you prefer to use the command-line, you use the `asr` tool to restore the image to the partition. Using `asr` requires the use of superuser or root privileges. The basic syntax is:

```
sudo asr restore -s <compressedimage> -t <targetvol> --erase
```

For example, restoring an image called “Installer.dmg” to the partition “ExtraHD” would be:

```
asr restore -s Installer.dmg -t ExtraHD --erase
```

For more information about `asr` and its capabilities, see the tool’s man page.

- ▶ **Tip:** You can use `asr` to restore a disk over a network, multicasting the blocks to client computers. Using the multicast server feature of `asr`, you could put a copy of the installer image on a partition of all computers that can receive the multicast packets. To successfully configure this, you’ll need the information in the tool’s man page.

The `asr` tool can also fetch the target image from an HTTP server using `http` or `https` URLs as its source, so the image doesn't need to reside on the target computer computer.

Step 4: Select the alternate partition as the startup disk.

After the partition is restored, it's a startup and installer disk for your server. You now need to start up the computer from that partition. After the computer is up and running, it is a Mac OS X Server installer, exactly as if you had started the computer from the DVD.

To start up the computer with the installation disc:

- 1 Turn on the computer and hold down the Option key.
- 2 Select the icon representing the installation partition and then click the right arrow.

You must use this method if you are starting up from an external DVD drive.

If you're installing on an Xserve, the procedure for starting up from a DVD may be different. For more information, see the *Xserve User's Guide* or *Quick Start* that came with your Xserve.

- 3 After the computer restarts, choose the language you want to use during installation, and click the arrow button.

The Installer is now running.

If you prefer to do this with the command-line, you can set the startup volume using the `systemsetup` tool. In version of Mac OS X Server since v10.4 or later, the `systemsetup` tool is at `/usr/sbin/systemsetup`.

If you are currently using the Mac OS X client during this process, the tool is at `/System/Library/CoreServices/RemoteManagement/ARDAgent.app/Contents/Support/systemsetup`.

You'll need to use the `-liststartupdisks`, and `-setstartupdisk` command options to find the newly restored installer volume, and select it as the startup disk. All commands issued with `systemsetup` must be run with superuser or root privileges.

The following is an example command to select the startup disk:

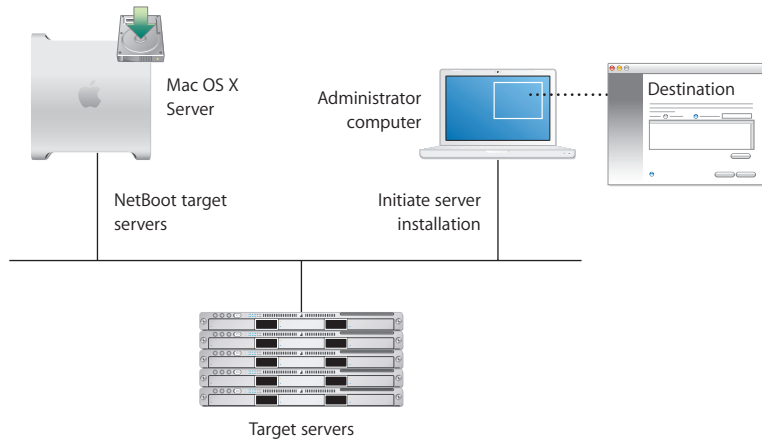
```
systemsetup -setstartupdisk "/Volumes/Mac OS X Server Install Disk"
```

Then issue the `shutdown -r` command to restart.

For more information about `systemsetup`, see *Command-Line Administration* and the tool's man page.

Starting Up from a NetBoot Environment

If you have an existing NetBoot infrastructure, this is the easiest way to perform mass installation and deployment. This method can be used for clusters that have no optical drive or existing system software, as shown in the following illustration:



It can also be used in environments where carrying large numbers of servers must be deployed in an efficient manner.

This section won't tell you how to create the necessary NetBoot infrastructure. If you want to set up NetBoot and NetInstall options for your network, servers, and client computers, see *System Imaging and Software Update Administration*.

This section has instructions to create a NetInstall image from the Mac OS X Server Install Disk, and start a server from it. There is no need to make preparations to the hard disk.

Step 1: Create a NetInstall image from the Install DVD

This step doesn't need to be done on the target computer. It can be done on an administrator computer that has enough free space to image the entire Install DVD.

- 1 Launch System Image Utility, in `/Applications/Server/`.
- 2 Select the Install DVD on the left, and choose NetInstall image on the right.
- 3 Click Continue.
- 4 Enter a name for the image, and a description.
This information is seen by clients selecting it a startup disk.
- 5 Click Create and then choose a save location for the disk image.

Upon completion, this image can be used with an existing NetBoot server to start up a server for installation.

For more information about NetInstall images and System Image Utility, including customization options, see *System Imaging and Software Update Administration*.

Step 2: Start up the computer from the NetBoot server

There are four ways of doing this, depending on your environment.

- In the target computer GUI, select the NetInstall disk from the Startup Disk pane of the System Preferences.
- Restart the computer, holding down the “n” key.
The first NetBoot server to respond to the computer will start up the computer with its default image.
- Restart the computer, holding down the Option key.
The computer will show you the available startup disks, locally on the computer and remotely from NetBoot and NetInstall servers. Select a disk and continue the startup.
- Use the command-line locally or remotely to specify the NetBoot server that the computer will start up from:

```
sudo bless --netboot --server bsdp://server.example.com
```

Preparing Disks for Installing Mac OS X Server

Before performing a clean installation of Mac OS X Server, you can partition the server computer's hard disk into multiple volumes, create a RAID set, or erase the target disk or partition.

If you're using an installation disc for Mac OS X Server v10.5 or later, you can perform these tasks from another networked computer using VNC viewer software, such as Apple Remote Desktop, before beginning a clean installation.

WARNING: Before partitioning a disk, creating a RAID set, or erasing a disk or partition on an existing server, preserve any user data you want to save by copying it to another disk or partition.

Choosing a File System

A file system is a method for storing and organizing computer files and the data they contain on a storage device such as a hard disk. Mac OS X Server supports several kinds of file systems to be used for hard disk storage. Each file system has its own strengths. You must decide which system fits your organization's needs.

For more information, see the following:

developer.apple.com/technotes/tn/tn1150.html

The following systems are available for use:

Mac OS Extended (Journaled) aka HFS+J

An HFS+J volume is the default file system for Mac OS X Server.

An HFS+J volume has an optional journal to speed recovery when mounting a volume that was not unmounted safely (for example, as the result of a power outage or crash). The journal makes it quick and easy to restore the volume structures to a consistent state, without having to scan all of the structures.

The journal is used only for the volume structures and metadata; it does not protect the contents of a fork. In other words, this journal protects the integrity of the underlying disk structures, but not any data that is corrupted due to a write failure or catastrophic power loss.

More information about HFS+J can be found in Apple's Developer Documentation at:

developer.apple.com/documentation/MacOSX/Conceptual/BPFileSystem/Articles/Comparisons.html

Mac OS Extended (Journaled, Case-Sensitive) aka HFSX

HFSX is an extension to HFS Plus and allows volumes to have case-sensitive file and directory names. Case-sensitive names means that you can have two objects whose names differ only by the case of the letters, in the same directory at the same time. For example, you could have Bob, BOB, and bob in the same directory as uniquely named files.

A case-sensitive volume is supported as a boot volume format. An HFSX file system for Mac OS X Server must be specifically selected when erasing a volume and preparing for initial installation. HFSX is an available format for the "erase and install" option for local installs. HFSX is *not* an available format for remotely controlled installations. If you are planning to use NFS, you should use case-sensitive HFSX.

An HFSX volume can be either case-sensitive or case-insensitive. Case sensitivity (or lack thereof) is global to the volume; the setting applies to all file and directory names on the volume. To determine whether an HFSX volume is case-sensitive, use the `keyCompareType` field of the B-tree header of the catalog file. A value of `kHFSBinaryCompare` means the volume is case-sensitive. A value of `kHFSCaseFolding` means the volume is case-insensitive.

Note: Do not assume that an HFSX volume is case-sensitive. Always use the `keyCompareType` to determine case sensitivity or case insensitivity. Additionally, don't assume your third-party software solutions work correctly with case sensitivity.

Important: Case-sensitive names do not ignore Unicode ignorable characters. This means that a single directory can have several names that would be considered equivalent using Unicode comparison rules, but they are considered distinct on a case-sensitive HFSX volume.

Partitioning a Hard Disk

Partitioning the hard disk creates a volume for server system software and one or more additional volumes for data and other software. Partitioning erases previous contents of the disk.

The minimum recommended size for an installation partition is 20 GB. A larger volume is recommended for a standard or workgroup configuration because they keep shared folders and group websites on the startup volume together with the server software.

Erasing a disk is another way of saying that you have given a disk a single volume partition and erased that volume.

Consider dedicating a hard disk or a volume of a partitioned hard disk to the server software. Put additional software, share points, websites, and so forth on other disks or volumes. With this approach, you can upgrade or reinstall the server software without affecting your other software or user data. If you must store additional software or data on the system volume, consider mirroring it to another drive.

- ▶ **Tip:** Having an extra, empty partition or two on the target installation disk can give you additional flexibility in installation and deployment. For example, additional space can give you a place to temporarily mirror your current installation before performing an in-place update, or it can give you a fast installer disk.

Partitioning a Disk Using Disk Utility

You can use the Installer to open the Disk Utility application and then use Disk Utility to erase the installation target volume or another volume. You can erase the target volume using the Mac OS Extended format, Mac OS Extended (Journaled) format, Mac OS Extended format (Case-Sensitive) format, and Mac OS Extended (Journaled, Case-Sensitive) format. You cannot partition the active startup disk or erase the active startup volume.

1 Launch Disk Utility.

If you are in the Installer, Disk Utility is available from the Utilities menu.

Otherwise, launch the application from `/Applications/Utilities/Disk Utility`.

2 Select the disk to be partitioned.

You can't select your current startup disk. Selecting a volume on the disk will allow you to erase the volume but will not create a different partition scheme.

3 Click Partition.

4 Choose your partition scheme and follow the instructions in the window to set all necessary parameters.

5 Click Apply.

You can find instructions for partitioning the hard disk into multiple volumes, creating a RAID set, and erasing the target disk or partition by viewing Disk Utility Help. To view Disk Utility Help, open Disk Utility on another Macintosh computer with Mac OS X v10.5 and choose Help > Disk Utility Help.

Partitioning a Disk Using the Command-line

You can use the `diskutil` command-line tool to partition and erase a hard disk. Normally, you would use a remote shell (SSH) to log in to the newly-started computer to use this method. The tool to partition disks is `diskutil`.

Just like using Disk Utility, you can erase the target volume using the Mac OS Extended format, Mac OS Extended (Journaled) format, Mac OS Extended format (Case-Sensitive) format, and Mac OS Extended (Journaled, Case-Sensitive) format.

- You cannot partition the active startup disk or erase the active startup volume.
- All potentially destructive `diskutil` operations must be done with superuser or root privileges.

Additional information about `diskutil` and other uses can be found in *Command-Line Administration*. For complete command syntax for `diskutil`, consult the tool's man page.

The specific command issued depends on your disk format needs and the hardware in use. Take care to use command-line arguments that apply to your specific needs.

The following command is a sample, which partitions a computer's only 120 GB hard disk into two equal 60 GB journaled HFS+ volumes ("BootDisk" and "DataStore"), which can start up a PowerPC-based Macintosh computer.

The basic syntax is:

```
diskutil partitionDisk device numberOfPartitions APMFormat <part1Format  
part1Name part1Size> <part2Format part2Name part2Size>
```

So the command is:

```
diskutil partitionDisk disk0 2 APMFormat JournaledHFS+ BootDisk 50%  
JournaledHFS+ DataStore 50%
```

Creating a RAID Set

If you're installing Mac OS X Server on a computer with multiple internal hard disks, you can create a Redundant Array of Independent Disks (RAID) set to optimize storage capacity, improve performance, and increase reliability in case of a disk failure.

For example, a mirrored RAID set increases reliability by writing your data to two or more disks at once. If one disk fails, your server uses one of the other disks in the RAID set.

You can use Disk Utility to set up a RAID set. There are two types of RAID sets and one additional disk option available in Disk Utility:

- **A striped RAID set (RAID 0)** splits files across the disks in the set. A striped RAID set improves the performance of your software because it can read and write on all disks in the set at the same time. You might use a striped RAID set if you are working with large files, such as digital video.
- **A mirrored RAID set (RAID 1)** duplicates files across the disks in the set. Because this scheme maintains two or more copies of the files, it provides a continuous backup of them. In addition, it can help keep data available if a disk in the set fails. Mirroring is recommended if you have shared files or applications that must be accessed frequently.

You can set up RAID mirroring after installing Mac OS X Server if you install on a disk that isn't partitioned. To prevent data loss, set up RAID mirroring as soon as possible.

- **A concatenated disk set** lets you use several disks as a single volume. This is not a true RAID set and offers no redundancy or performance increase.

You can combine RAID sets to combine their benefits. For example, you can create a RAID set that combines the fast disk access of a striped RAID set and the data protection of a mirrored RAID set. To do this, create two RAID sets of one type and then create a RAID set of another type, using the first two RAID sets as the disks.

The RAID sets you combine must be created with Disk Utility or `diskutil` in Mac OS X v10.4 or later.

You cannot mix the method of partitioning used on the disks in a RAID set. (The PPC platform is APMFormat and the Intel platform is GPTFormat.)

Mac Pro desktop computers and Intel-based Xserves can boot from a software RAID volume. Some Intel-based Macs do not support booting from software RAID volumes. If you try to start these Intel-based Macs from a software RAID volume, the computer might start up with a flashing question mark.

The following computers do not support booting from software RAID volumes:

- iMac (Early 2006)
- Mac mini (Early 2006)

No PPC-based Macs support booting from software RAID volumes.

If you need more sophisticated RAID support, consider a hardware RAID. It has specially dedicated RAID hardware and can contain over 5 terabytes of storage.

Creating a RAID Set Using Disk Utility

You can use the Installer to open Disk Utility and then use Disk Utility to create the RAID set from available disks. Creating a RAID set erases the contents of the disks involved, so it isn't necessary to erase the disks before creating the RAID set.

The RAID set volumes can be Mac OS Extended format, Mac OS Extended (Journaled) format, Mac OS Extended format (Case-Sensitive) format, Mac OS Extended (Journaled, Case-Sensitive) format, and MS-DOS FAT format. For more information about volume formats, see “Preparing Disks for Installing Mac OS X Server” on page 91.

You cannot create a RAID set from the startup disk.

To create a RAID set using Disk Utility:

- 1 Launch Disk Utility.

If you are in the Installer, Disk Utility is Available from the Utilities menu; otherwise, launch the application from /Applications/Utilities/Disk Utility.

- 2 Select the disk to be part of the RAID set.

You can't select your startup disk.

When creating RAID sets or adding disks, specify the entire disk instead of a partition on that disk.

- 3 Click RAID.
- 4 Choose your RAID set type.
- 5 Drag the disks to the window.
- 6 Follow the instructions in the window to set parameters.
- 7 Click Create.

You can find instructions for partitioning the hard disk into multiple volumes, creating a RAID set, and erasing the target disk or partition by viewing Disk Utility Help. To view Disk Utility Help, open Disk Utility on another Macintosh computer with Mac OS X v10.5 and choose Help > Disk Utility Help.

Creating a RAID Set Using the Command-line

You can use the `diskutil` command-line tool to create a RAID set. Normally, you would use a remote shell (SSH) to log in to the newly-started computer to use this method.

You can use `diskutil` to create a RAID volume that is Mac OS Extended format, Mac OS Extended (Journaled) format, Mac OS Extended format (Case-Sensitive) format, Mac OS Extended (Journaled, Case-Sensitive) format, or MS-DOS FAT format. However keep in mind the following:

- You cannot create a RAID from the startup disk.
- When creating RAID sets or adding disks, specify the entire disk instead of a partition on that disk.
- All potentially destructive `diskutil` operations must be done with superuser or root privileges.

Additional information about `diskutil` and other uses can be found in *Command-Line Administration*. For complete command syntax for `diskutil`, consult the tool's man page.

Use command-line arguments that apply to your specific needs. The following command is a sample, which creates a single mirrored RAID set (RAID 1) from the first two disks installed in the computer (disk0 and disk1), with the resulting RAID volume called MirrorData.

The basic syntax is:

```
diskutil createRAID mirror setName format device device ...
```

So the command is:

```
diskutil createRAID mirror MirrorData JournaledHFS+ disk0 disk1
```

Erasing a Disk or Partition

You have several options for erasing a disk, depending on your preferred tools and your computing environment:

- **Erasing a Disk Using the Installer:** You can erase a disk or partition while using the Mac OS X Server Installer. When you select the target volume in the Installer, you can also select an option to have the target disk or partition erased during installation using the Mac OS Extended (Journaled) format. This is the most recommended format for a Mac OS X Server startup volume.
- **Erasing a Disk Using Disk Utility:** You can use the Installer to open the Disk Utility application and then use it to erase the target volume or another volume. You can erase the target volume using the Mac OS Extended format or Mac OS Extended (Journaled) format. You can erase other volumes using either of those formats, Mac OS Extended format (Case-Sensitive) format, or Mac OS Extended (Journaled, Case-Sensitive) format.

You can erase but not partition a disk or partition while using the Mac OS X Server Installer. When you select the target volume in the Installer, you can also select an option to have the target disk or partition erased during installation using the Mac OS Extended (Journaled) format. This is the recommended format for a Mac OS X Server startup volume.

You can find instructions for partitioning the hard disk into multiple volumes, creating a RAID set, and erasing the target disk or partition by viewing Disk Utility Help. To view Disk Utility Help, open Disk Utility on another Macintosh computer with Mac OS X v10.5 and choose Help > Disk Utility Help.

- **Erasing a Disk Using the Command-line:** Finally you can use the command-line to erase disks using the tool diskutil. Erasing a disk using diskutil results in losing all of the volume partitions. The command to erase a complete disk is:

```
diskutil eraseDisk format name [OS9Drivers | APMFormat | MBRFormat |  
GPTFormat] device
```

For example:

```
diskutil eraseDisk JournaledHFS+ MacProHD GPTFormat disk0
```

There is also an option to securely delete data by overwriting the disk with random data multiple times. For more details, see `diskutil`'s man page.

To erase a single volume on a disk, a slightly different command is used:

```
diskutil eraseVolume format name device
```

For example:

```
diskutil eraseVolume JournaledHFS+ UntitledPartition /Volumes/  
OriginalPartition
```

Additional information about `diskutil` and other uses can be found in *Command-Line Administration*. For complete command syntax for `diskutil`, consult the tool's man page.

Identifying Remote Servers When Installing Mac OS X Server

For remote server installations, you need to know this information about the target server:

- **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 you must enter the IP address of the server (in IPv4 format: 000.000.000.000) if it resides on a different subnet. Information provided for servers in the list includes IP address, host name, and Media Access Control (MAC) address (also called hardware or Ethernet address).

If you use VNC viewer software to remotely control installation of Mac OS X Server v10.5 or later, it may let you select the target server from a list of available VNC servers. If not, you must enter the IP address of the server (in IPv4 format: 000.000.000.000).

The target server's IP address is assigned by a DHCP server on the network. If no DHCP server exists, the target server uses a 169.xxx.xxx.xxx address unique among servers on the local subnet. Later, when you set up the server, you can change the IP address.

If you don't know the IP address and the remote server is on the local subnet, you can use the `sa_srchr` command to identify computers on the local subnet where you can install server software. Enter the following from an existing computer with Mac OS X Server Tools installed:

```
/System/Library/Serverssetup/sa_srchr 224.0.0.1
```

This command will return the IP address, and the EthernetID (in addition to other information) of servers on the local subnet which have started up from the installation disk.

- **The preset password for the target server:** The password consists of the first eight digits of the server's built-in hardware serial number. To find a server's serial number, look for a label on the server. Older computers have no built-in hardware serial numbers; for these systems, use 12345678.

Installing Server Software Interactively

You can use the installation disc to install server software interactively on a local server, on a remote server, or on a computer with Mac OS X pre-installed.

Installing Locally from the Installation Disc

You can install Mac OS X Server directly onto a computer with a display, a keyboard, and an optical drive attached, as shown in the following illustration:



If you have an Install DVD, the optical drive must be able to read DVD discs.

You can also install directly onto a computer that lacks a display, keyboard, and optical drive capable of reading your installation disc. In this case, you start the target computer in target disk mode and connect it to an administrator computer using a FireWire cable.

You use the administrator computer to install the server software on the target computer's disk or partition, which appears as a disk icon on the administrator computer.

These instructions assume you have started up the computer using the Install DVD, installer partition, or NetInstall disk. If you have not, see the relevant instructions beginning at "About Starting Up for Installation" on page 83.

To install server software locally:

- 1 After the computer starts, choose the language you want the server to use and click Continue.
- 2 When the Installer opens, if you want to perform a clean installation, optionally use the Utilities menu to open Disk Utility to prepare the target disk or partition before proceeding.

If you have not previously prepared your disk for installation, you can do so now with Disk Utility. For more instructions on preparing your disk for installation, see "Preparing Disks for Installing Mac OS X Server" on page 91.

- 3 Proceed through the Installer's panes by following the onscreen instructions.

- 4 When the Select a Destination pane appears, select a target disk or volume (partition) and make sure it's in the expected state.

If you're doing a clean installation, you can click Options to format the destination disk or volume in Mac OS Extended (Journaled) format. Select Erase to format the disk in Mac OS Extended (Journaled) format; then click OK.

If the volume you selected contains Mac OS X Server v10.3.9 or 10.2.8 and you want to upgrade, click Options, select "Don't erase," and then click OK.

Important: When you perform an upgrade, make sure that saved setup data won't be inadvertently detected and used by the server. If saved setup data is used, the server settings are not compatible with the saved settings and can cause unintended consequences. For more information, see "How a Server Searches for Saved Setup Data" on page 124.

- 5 Proceed through the Installer's panes by following the onscreen instructions.

After installation is complete, the computer restarts and you can perform initial server setup.

- 6 If you're using an administrator computer to install onto a server that's in target disk mode and connected using a FireWire cable:

- a Quit Server Assistant when it starts automatically on the administrator computer.

- b Shut down the administrator computer and the server.

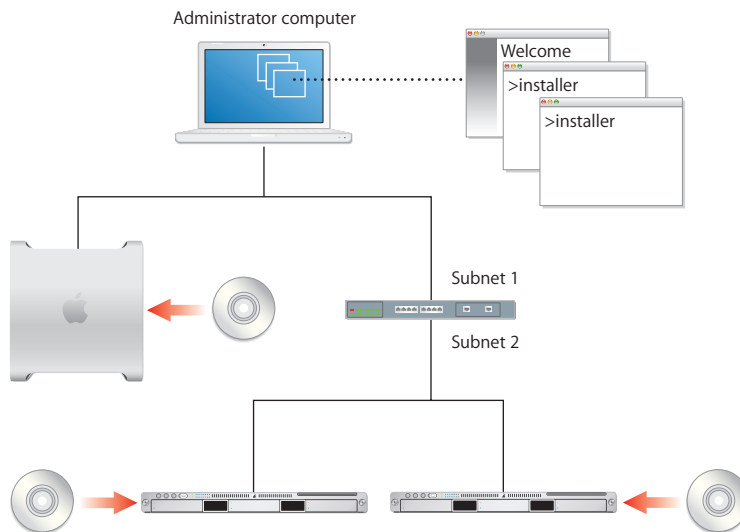
- c Start up the administrator computer and the server normally (not in target disk mode).

Now you can use Server Assistant from the administrator computer to remotely set up the server.

Chapter 6, "Initial Server Setup," on page 107 describes how to set up a server locally or remotely.

Installing Remotely with Server Assistant

To install Mac OS X Server on a remote server from the server Install DVD, installation partition, or NetInstall disk, you need an administrator computer from which to use Server Assistant to manage the installation:



After the computer starts up, you can control and manage other servers from an administration computer.

Important: If you have administrative applications and tools from Mac OS X Server v10.4 Tiger or earlier, do not use them with Leopard Server.

To use the Installer user interface, use VNC to view and interact with the remote installer. For more information, see “Installing Remotely with VNC” on page 102.

These instructions assume you have started up the computer using the Install DVD, installer partition, or NetInstall disk. If you have not, see the instructions at “About Starting Up for Installation” on page 83.

You don’t need to be an administrator on the local computer to use Server Assistant.

To install on a remote server by using Server Assistant:

- 1 After the target computer has started from the server Install DVD, installation partition, or NetInstall disk, launch Server Assistant in the /Applications/Server/ folder on the administrator computer.
- 2 Select “Install software on a remote server.”
- 3 For every intended target server, identify the target server and add it to the list.

If it's on the local subnet, select it in the list; otherwise, click the Add (+) button and enter an IP address in IPv4 format (000.000.000.000).

If you already have a saved server list, load it now by selecting File > Load Server List.

- 4 When prompted for a password, enter the first eight digits of the server's built-in hardware serial number.

To find a server's serial number, 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.

- 5 After you finish adding servers to the list, save this list for future use by selecting File > Save Server List.
- 6 Proceed by following the onscreen instructions.
- 7 When the Volumes pane appears, select a target disk or volume (partition), make sure it's in the expected state, and click Continue.

If the volume you selected contains Mac OS X Server v10.4.10 or v10.3.9 and you want to upgrade, select "Don't erase"; otherwise, select Erase to format the disk in Mac OS Extended (Journaled) format; then click OK.

WARNING: When you perform an upgrade, make sure that saved setup data won't be detected and used by the server. If saved setup data is used, the server settings are not compatible with the saved settings and can cause unintended consequences. For more information, see "How a Server Searches for Saved Setup Data" on page 124.

- 8 Proceed by following the onscreen instructions.

While installation proceeds, you can open another Server Assistant window to install server software on other computers. Choose File > New Window to do so.

After installation is complete, the target server restarts and you can perform initial server setup. Chapter 6, "Initial Server Setup," on page 107 describes how.

Installing Remotely with VNC

If you're using an installation disc for Mac OS X Server v10.5 or later, you can control installation from another computer using open source VNC viewer software or Apple Remote Desktop. This allows you to remotely control preparation of the target disk or partition before beginning installation.

You can partition the hard disk into multiple volumes, create a RAID set, or erase the target disk or partition.

The process for remotely installing with VNC is the same as installing locally at the keyboard and monitor, except that you must first connect to the VNC server on the target computer with a VNC client, like Apple Remote Desktop.

For information about connecting to a computer running from an Install DVD, see “Remotely Accessing the Install DVD” on page 84.

For information about running the installer locally, see “Installing Locally from the Installation Disc” on page 99.

Using the installer Command-Line Tool to Install Server Software

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

- See *Command-Line Administration*.
- Open the Terminal application and type `installer`, `installer -help`, or `man installer`.

These instructions assume you have started up the computer using the Install DVD, installer partition, or NetInstall disk. If you have not, see the relevant instructions beginning at “About Starting Up for Installation” on page 83.

To use installer to install server software:

- 1 Start a command-line session with the target server by choosing from the following:
 - Installing a local server: When the Installer opens choose Utilities > Open Terminal to open the Terminal application.
 - Installing a remote server: From Terminal on an administrator computer or from a UNIX workstation, establish an SSH session as the root user with the target server, substituting the target server’s actual IP address for `<ip address>`:

```
ssh root@<ip address>
```

- ▶ If you don’t know the IP address and the remote server is on the local subnet, you can use the `sa_srchr` command to identify computers on the local subnet where you can install server software:

```
/System/Library/Serverssetup/sa_srchr 224.0.0.1  
mycomputer.example.com#PowerMac4,4#<ip address>#<mac address>#Mac OS X  
Server 10.5#RDY4PkgInstall#2.0#512
```

You can also use Server Assistant to generate information for computers on the local subnet. Open Server Assistant, select “Install software on a remote computer,” and click Continue to access the Destination pane and generate a list of servers awaiting installation.

- 2 When prompted for a password, enter the first eight digits of the server’s built-in hardware serial number.

To find a server's serial number, look for a label on the server. If the target computer had been set up as a server, you'll also find the hardware serial number in `/System/Library/Serverssetup/SerialNumber`.

If you're installing on an older computer that has no built-in hardware serial number, use 12345678 for the password.

3 Identify the target-server volume where you want to install the server software.

To list the volumes available for server software installation from the installation disc, type this command:

```
/usr/sbin/installer -volinfo -pkg /System/Installation/Packages/  
OSInstall.mpkg
```

You can also identify a NetInstall image you've created and mounted:

```
/usr/sbin/installer -volinfo -pkg /Volumes/ServerNetworkImage10.5/  
System/Installation/Packages/OSInstall.mpkg
```

The list displayed reflects your particular environment, but here's an example showing three available volumes:

```
/Volumes/Mount 01  
/Volumes/Mount1  
/Volumes/Mount02
```

4 If you haven't already done so, prepare the disks for installation.

For more information about preparing the disks for installation, see "Preparing Disks for Installing Mac OS X Server" on page 91.

If the target volume has Mac OS X Server v10.4.10 or 10.3.9 installed, when you run `installer` it will upgrade the server to v10.5 and preserve user files.

If you're not upgrading but performing a clean installation, back up the user files you want to preserve, then use `diskutil` to erase the volume and format it to enable journaling:

```
/usr/sbin/diskutil eraseVolume HFS+ "Mount 01" "/Volumes/Mount 01"  
/usr/sbin/diskutil enableJournal "/Volumes/Mount 01"
```

You can also use `diskutil` to partition the volume and to set up mirroring. For more information about the command, 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 if you need to reinstall or upgrade system software. If you must store additional software or data on the system partition, consider mirroring the drive.

5 Install the operating system on a volume from the list generated in step 3.

For example, to use Mount 01 in the example in step 4 to install from a server installation disc, enter:


```
/usr/sbin/installer -verboseR -lang en -pkg /System/Installation/  
Packages/OSInstall.mpkg -target "/Volumes/Mount 01"
```

If you're using a NetInstall image, the command identifies them as step 3 shows.

When you enter the `-lang` parameter, use one of the following values: `en` (for English), `de` (for German), `fr` (for French), or `ja` (for Japanese).

During installation, progress information is displayed. While installation proceeds, you can open another Terminal window to install server software on another computer.

- 6 When installation from the disc is complete, restart the server by entering:

```
/sbin/reboot
```

or

```
/sbin/shutdown -r
```

Server Assistant opens when installation is complete. You can now proceed to set up the server. For more information, see “Initial Server Setup” on page 107.

Installing Multiple Servers

You can use Server Assistant, VNC viewer software, or the `installer` tool to initiate multiple server software installations. After using Server Assistant to initiate server software installation on more than one remote computer, you can choose File > New Window to install the software on another batch of computers.

When running Server Assistant from an administration computer to install on multiple machines, group the same hardware configurations together. For example, choose all Intel Xserve machines or all G4 Mac minis.

After using a VNC viewer to control installation of Mac OS X Server v10.5 or later on one remote computer, you can use the VNC viewer to open a connection to another remote computer and control installation on it. Because this involves interacting with each server individually, it is a less efficient method of installing on multiple servers.

The most efficient method of installation would be completely automated. Opening the Terminal application and using the `installer` tool to initiate each server software installation doesn't accomplish this efficiently. However, scripting the command-line tool (using known values for server IP addresses, for example) to automate multiple simultaneous installations can be very efficient. To completely automate server installation, you must script the `installer` tool and have a high measure of control over the network infrastructure.

For example, to have known IP addresses and the appropriate hardware serial numbers included in your script, you cannot rely on the randomly assigned IP addresses. You can use DHCP assigned static addresses to remove that uncertainty and ease your scripting considerations.

The methods, scripting languages, and possibilities are too many to list in this guide.

Upgrading a Computer from Mac OS X to Mac OS X Server

You can use the Install DVD for Mac OS X Server v10.5 to upgrade a desktop computer that has the following characteristics:

- Has Mac OS X v10.5 or later installed
- Has an Intel processor
- Was introduced in summer 2006 or later
- Meets the system requirements in “System Requirements for Installing Mac OS X Server” on page 81

To upgrade a computer from Mac OS X to Mac OS X Server:

- 1 Start up the computer from the hard disk, as you would for normal use.

Do not use an installation disc.

- 2 Insert the Install DVD, open the Other Installs folder, and double-click MacOSXServerInstall.mpkg to run the Installer.

When the Installer finishes, your computer restarts automatic[ally and Server Assistant opens to let you set up the server.

- 3 After the server restarts, use Software Update to install server software updates.

How to Keep Current

After you’ve set up your server, you’ll want to update it when Apple releases server software updates.

There are several ways to access update releases of Mac OS X Server:

- In Server Admin, select a server in the Servers list, then click the Server Updates button.
- Use the Software Update pane of System Preferences.
- Use the `softwareupdate` command-line tool.
- Use the server’s software update service.
- Download a disk image of the software update from:
www.apple.com/support/downloads

Basic characteristics of your Mac OS X Server are established during server setup. The server can operate in three different configurations: advanced, standard, and workgroup. These instructions assume you have chosen the advanced configuration.

After installing server software, the next task is to set up the server. There are several ways to set up a server:

- Set up one or more servers interactively.
- Automate the setup by using setup data you've saved in a file or in a directory the servers are configured to access.

Information You Need

To understand and record information for each server you want to set up, see the “Mac OS X Server Advanced Worksheet” in the appendix on page 197. The following information provides supplemental explanations for some items on the worksheet.

When you're upgrading from Mac OS X Server v10.4.10 or v10.3.9, Server Assistant displays the existing server settings, but you can change them. Use the “Mac OS X Server Advanced Worksheet” to record settings you want the v10.5 server to use.

Postponing Server Setup Following Installation

Server Assistant opens automatically on a server that hasn't been set up, and waits for you to begin the setup process. To set up the server later, you can postpone the setup process by using the server's keyboard, mouse, and display.

To postpone setting up Mac OS X Server:

- In Server Assistant, press Command-Q on the server's keyboard, and then click Shut Down.

When you restart the server, Server Assistant opens again.

If you're setting up a server without a keyboard or display, you can enter commands in the Terminal application to shut down the server remotely. For information about using the command-line to connect to a remote computer and shut it down, see *Command-Line Administration*.

Connecting to the Network During Initial Server Setup

Try to place a server in its final network location (subnet) before setting it up for the first time. If you're concerned about preventing unauthorized or premature access during setup, you can set up a firewall to protect the server while you're finalizing its configuration.

If you can't avoid moving a server after initial setup, you must change settings that are sensitive to network location before it can be used. For example, the server's IP address and host name, stored in directories and configuration files on the server, must be updated. For more information, see "Changing the Server's Host Name After Setup" on page 146.

Configuring Servers with Multiple Ethernet Ports

Your server has a built-in Ethernet port and might have additional Ethernet ports built in or added on.

When you're using Server Assistant to interactively set up servers, all of a server's available Ethernet ports are listed and you select them to activate and configure. When you work in Server Assistant's offline mode, you click an Add button to create a list of ports to configure.

If you enable more than one port, you specify the order for the ports to be used by the server when routing traffic to the network. Although the server receives network traffic on any active port, network traffic initiated by the server is routed through the first active port.

For a description of port configuration attributes, see the "Mac OS X Server Advanced Worksheet" on page 197.

About Settings Established During Initial Server Setup

During server setup, the following basic server settings are established:

- The language to use for server administration and the computer keyboard layout is defined.
- The server software serial number is set.
- A server administrator user is defined and the user's home folder is created.
- Default AFP and FTP share points, such as Shared Items, Users, and Groups, are defined.
- Basic Open Directory information is set up. At a minimum, a local directory domain is created. You can also set up an LDAP directory for other computers to use or configure the server to obtain directory information from other servers.
- The server's host name, computer name, and local hostname are set. You can specify the computer name and local hostname, but Server Assistant sets the host name to AUTOMATIC in `/etc/hostconfig`. This setting causes the server's host name to be the first name that's true in this list:
 - The name provided by the DHCP or BootP server for the primary IP address
 - The first name returned by a reverse DNS (address-to-name) query for the primary IP address
 - The local hostname
 - The name "localhost"
- Network interfaces (ports) are configured. TCP/IP and Ethernet settings are defined for each port you want to activate.
- Network time service can be set up.

If you're upgrading, the current basic settings are displayed during the setup process, but you can change them. Other settings, such as share points you've defined and services you've configured, are preserved. For a complete description of what's upgraded and actions, see *Upgrading and Migrating*.

You can perform initial server setup only once without reinstalling a server. To change settings established during setup, you have alternative means to do so. For example, you can use Server Admin or Directory Utility to manage Open Directory settings.

Specifying Initial Open Directory Usage

During setup of an advanced configuration, you specify how the server initially stores and accesses user accounts and other directory information. You choose whether the server connects to a directory system or works as a standalone server.

After setup, you can create or change a connection to a directory system by using Directory Utility, or you can make the server an Open Directory master or replica by using Server Admin to change the server's Open Directory service settings. For information about changing directory services, see *Open Directory Administration*.

When you set up a server initially, you specify its directory services configuration. Choices are:

- **No change**, available only when upgrading from Mac OS X Server v10.4.10 or 10.3.9.
- **Standalone Server**, used to set up only a local directory domain on the server.
- **Connected to a Directory System**, used to set up the server to obtain directory information from a shared directory domain that's already been set up on another server.

In all these cases, Open Directory authentication is set up on the server and used by default for any new users added to domains that reside on the server.

If you're setting up multiple servers and one or more of them will host a shared directory, set up those servers before setting up servers that will use those shared directories.

Note: If you connect Mac OS X Server v10.5 to a directory domain of Mac OS X Server v10.2 or earlier, users defined in the older directory domain cannot be authenticated with the MS-CHAPv2 method.

This method may be required to securely authenticate users for the VPN service of Mac OS X Server v10.5. Open Directory in Mac OS X Server v10.5 supports MSCHAPv2 authentication, but Password Server in Mac OS X Server v10.2 doesn't support MS-CHAPv2.

After setup, use the Directory Utility or Server Admin applications to refine the server's directory configuration, if necessary. Directory Utility lets you set up connections to multiple directories, including Active Directory and other non-Apple directory systems, and specify a search policy (the order in which the server should search through the domains). Server Admin lets you set up replicas of an Open Directory master and manage other aspects of a server's directory service configuration.

Open Directory Administration can help you decide which directory usage setup option is right for you. If you're upgrading, the best choice is usually "No change." If you're setting up a new server, the simplest choice is "Standalone Server." After initial server setup, you can use Directory Utility or Server Admin to adjust and finalize the directory setup.

Not Changing Directory Usage When Upgrading

When setting up a server that you're upgrading to v10.5 from v10.3.9 or 10.2.8, and you want the server to use the same directory setup it's been using, choose "No change" in the Directory Usage pane in Server Assistant.

Even when you want to change the server's directory setup, selecting "No change" is the safest option, especially if you're considering changing a server's shared directory configuration. Changing from hosting a directory to using another server's shared directory or vice versa, or migrating a shared NetInfo domain to LDAP are examples of directory usage changes you should make *after* server setup in order to preserve access to directory information about your network.

For information about all the directory usage options available to you and how to use Directory Utility and Server Admin to make directory changes, see *Open Directory Administration*. For information about how to continue using existing directory data when you change directory service settings, see *Upgrading and Migrating*.

If you choose the "No change" option and the server wasn't using a Password Server, Open Directory authentication is set up. When you add users to any Apple directory domain residing on the server, their passwords are validated by default using Open Directory authentication.

Setting Up a Server as a Standalone Server

A standalone server stores and accesses account information in its local directory domain. The standalone server uses its local directory domain to authenticate clients for its file, mail, and other services. Other servers and client computers can't access the standalone server's local directory domain.

Open Directory authentication is also set up on the server. By default, Open Directory authentication is used when a user is added to the local domain.

When a user attempts to log in to the server or use one of its services that require authentication, the server authenticates the user by consulting the local database. If the user has an account on the system and supplies the appropriate password, authentication succeeds.

Setting Up a Server to Connect to a Directory System

If it's connected to another directory system, your server stores and accesses account information in another server's shared directory and can use the other directory system to authenticate clients for file, mail, and other services. Your server can also use its local directory domain for accounts and authentication.

You can integrate your server with a variety of directory systems by choosing one of the following options during setup:

- **Open Directory Server:** Your server can store and access directory information about an Open Directory server using LDAP. With this option, you need to know the DNS name or IP address of the Open Directory server.
- **As Specified by DHCP Server:** Your server will obtain information for connecting to a directory system from a DHCP server. The DHCP server must be set up to provide the address and search base of an LDAP server (DHCP option 95). The directory service and DHCP service are independent. They don't need to be provided by the same server.
- **Other Directory Server:** If you need to integrate the server with another kind of directory system or with multiple directory systems, choose this option and set up the connections later using the Directory Utility application.

This option lets you integrate your server into almost any existing directory service, including Microsoft Active Directory, Novell eDirectory, another non-Apple directory, or an NIS domain. For information about using Directory Utility, see *Open Directory Administration* or open Directory Utility and then use the Help menu.

If you set up your server to connect to an Open Directory server that has Mac OS X Server v10.3 or earlier, you may not be able to take advantage of some features:

- VPN service requires MS-CHAP2 authentication, which isn't available in v10.2 or earlier.
- Replication isn't supported by v10.2 or earlier.
- Kerberos configuration is much more complex in v10.2. In addition, automatic synchronization of Kerberos and Password Server requires v10.3 or later.
- In v10.3 and earlier, trusted directory binding, support for LDAP subdomains, and Directory Utility controls aren't available.

Using Interactive Server Setup


The simplest way to set up a small number of servers is to use Server Assistant's guided interview process after establishing a connection with each server in turn. You provide server setup data interactively, then initiate setup immediately. If you have only a few servers to set up, the interactive approach is useful. You can use the interactive approach to set up a local server, a remote server, or several remote servers.

To use this approach, open Server Assistant, connect to one or more target servers, supply setup data, and then initiate the setup immediately.

This is the technique you use to set up a local server, as "Setting Up a Local Server Interactively" on page 113 describes. You can also use this interactive approach to set up a remote server from an administrator computer. For instructions, see "Setting Up a Remote Server Interactively" on page 114.

When multiple remote servers can use the same setup data, you can supply the data and then initiate setup of all the servers at once, using a batch approach. When running Server Assistant from an administration computer to set up multiple servers, group the same hardware configurations together. For example, choose all Intel Xserve machines or all G4 Mac minis.

This technique, shown on the left side of the following illustration, requires that network identifiers for all target servers be set using DHCP or BootP. For instructions, see “Setting Up Multiple Remote Servers Interactively in a Batch” on page 115.



To customize the setup of individual servers, you can manage each setup individually from a different Server Assistant window. This approach is shown on the right side of the illustration above. For instructions, see “Setting Up a Remote Server Interactively” on page 114.

Although the previous illustration shows target servers on the same subnet as the administrator computer in one scenario and target servers on a different subnet in the other scenario, both setup scenarios can be used to set up servers on the same and different subnets.

If a target server is on a different subnet, you must supply its IP address. Servers on the same subnet are listed by Server Assistant, so you select one or more servers in the list.

Setting Up a Local Server Interactively

After you install server software, you can use the interactive approach to set it up locally if you have physical access to the computer.

This setup assumes you are using the Advanced server configuration mode. Don't try to use these instructions with Standard or Workgroup modes.

To set up a local server interactively:

- 1 Fill out the “Mac OS X Server Advanced Worksheet” on page 197.
Supplemental information appears in “Information You Need” on page 107.
When the server restarts, Server Assistant opens.
- 2 Enter the setup data you’ve recorded on the “Mac OS X Server Advanced Worksheet” as you move through the Assistant’s panes, following the onscreen instructions.
Make sure that DHCP or DNS servers you specify for the server you’re setting up to use are running.
After you enter setup data, Server Assistant displays a summary of the data.
- 3 Review the setup data you entered and, if necessary, click Go Back to change it.
- 4 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 a configuration file or directory record, select “Save in Encrypted Format” and then enter and verify a passphrase. You must supply the passphrase before a target server can use an encrypted setup file.
- 5 To initiate setup of the local server, click Apply.
- 6 When server setup is complete, click Restart Now.
Now you can log in as the server administrator user created during setup to configure services.

Setting Up a Remote Server Interactively

After server software is 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 a remote server interactively:

- 1 Fill out the “Mac OS X Server Advanced Worksheet” on page 197.
Supplemental information appears in “Information You Need” on page 107.
- 2 Make sure the target server is running.
- 3 On an administrator computer, open Server Assistant in /Applications/Server/.
You don’t need to be an administrator on the administrator computer to use Server Assistant.
- 4 In the Welcome pane, select “Set up a remote server” and click Continue.
- 5 In the Destination pane, put a check in the Apply column for the remote server you want to set up, enter 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.

- 6 For the server configuration type, select “Advanced.”
- 7 In the Language pane, specify the language you want to use to administer the target server.
- 8 If you are using saved setup data, do the following:
 - a In the Language pane, choose File > Open Configuration File or File > Open Directory Record to load the saved setup data you want to use. If the saved setup data is encrypted, enter the passphrase when prompted.
 - b Optionally choose View > Jump to Review to review the setup data, then use Go Back as necessary to change it.
- 9 If you are entering setup data, do the following:
 - a Click Continue and enter the setup data as you move through the Assistant’s panes, following the onscreen instructions, and click Continue.
 - b Make sure that DHCP or DNS servers you specify for the server you’re setting up are running.
- 10 After you specify setup data, review the summary displayed by Server Assistant and optionally click Go Back to change data.
- 11 To save the setup data as a text file or in a form you can use for automatic server setup (as a saved setup file or saved directory record), click Save As.
To encrypt a configuration file or directory record, select “Save in Encrypted Format” and then enter and verify a passphrase.
You must supply the passphrase before a target server can use an encrypted setup file.
- 12 To initiate setup of the remote target server, click Apply.
- 13 When server setup is complete, click Continue Now.
The target server restarts and you can log in as the server administrator user you created during setup to configure services.

Setting Up Multiple Remote Servers Interactively in a Batch

You can use the interactive approach to set up multiple servers as a batch if:

- The servers are accessible from an administrator computer
- The servers use the same chip platform (for example, Intel-based or PowerPC-based)
- The servers use the same setup data, except for server software serial numbers and network identities (host name, computer name, and local hostname)
- Network identities are provided by a DHCP or BootP server

When running Server Assistant from an administration computer to set up multiple servers, group the same hardware configurations together. For example, choose Intel Xserve machines or G4 Mac minis.

If you have servers with different configuration files, you can open a Server Assistant window for each server type. This way you can group servers by platform, settings, subnet, or other criteria you choose.

To set up multiple remote servers interactively in a batch:

- 1 Fill out the “Mac OS X Server Advanced Worksheet” on page 197 with the settings you want to use for all servers you want to set up.

Supplemental information appears in “Information You Need” on page 107. The preface tells you where else you can find the “Mac OS X Server Advanced Worksheet”.

- 2 Make sure the target servers and any DHCP or DNS servers you want them to use are running.
- 3 On an administrator computer that can connect to all the target servers, open Server Assistant. It’s located in /Applications/Server/. You don’t have to be an administrator on the administrator computer to use Server Assistant.
- 4 In the Welcome pane, select “Set up a remote server” and click Continue.
- 5 In the Destination pane, put a check in the Apply column for each remote server you want to set up. Then enter the preset password in the Password field for each server, and click Continue to connect to the servers.

If you don’t see a target server you want to set up on the list, click Add to add it.

- 6 In the Language pane, specify the language you want to use to administer the target servers.
- 7 If you are using saved setup data, do the following:

In the Language pane, choose File > Open Configuration File or File > Open Directory Record to load the saved setup data you want to use. If the saved setup data is encrypted, enter the passphrase when prompted.

Optionally choose View > Jump to Review to review the setup data, then use Go Back as necessary to change it.

- 8 If you are entering setup data, do the following:

Click Continue and enter the setup data as you move through the Assistant’s panes, following the onscreen instructions, and click Continue.

Make sure that any DHCP or DNS servers you specify for the server you’re setting up to use are running.

- 9 After all setup data is specified, review the summary displayed by Server Assistant and optionally click Go Back to change data.
- 10 To save the setup data as a text file or in a form you can use for automatic server setup (as a saved setup file or saved directory record), click Save As.

To encrypt a configuration file or directory 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.

- 11 To initiate server setup, click Apply.
- 12 To initiate setup of the remote target server, click Apply.
- 13 When server setup is complete, click Continue Now.

The target servers restart and you can log in as the server administrator user created during setup to configure their services.

Using Automatic Server Setup

When you have more than a few servers to set up, consider using automatic server setup. This approach also provides a way to preserve setup data so it can be reused if you need to reinstall server software.

The automatic approach is useful when you:

- Have more than a few servers to set up
- Want to prepare for setting up servers that aren't yet available
- Want to save setup data for backup purposes
- Need to reinstall servers frequently

To use automatic server setup, you use Server Assistant to specify setup data for each computer or batch of computers; then you save the data in a file or in a directory to create setup data, as shown in the following illustration:



Finally, you provide that setup data to the target servers. You can provide the data using a variety of methods, like storing files on the hard disk, removable storage, or saving it in directory entries. By default, saved setup data is encrypted for extra security.

When a server starts up for the first time, it searches for automatic setup data to configure itself before it starts the interactive Setup Assistant.

Automatic server setup requires two main steps:

Step 1: Create the setup data files

The following sections can help you create setup data files.

- “Setup Data Saved in a File” on page 119
- “Setup Data Saved in a Directory” on page 120
- “Creating and Saving Setup Data” on page 118
- “Keeping Backup Copies of Saved Setup Data” on page 121

Step 2: Make the setup data files available to a freshly installed server

The following sections can help you make the data available to the servers:

- “How a Server Searches for Saved Setup Data” on page 124
- “Setting Up Servers Automatically Using Data Saved in a File” on page 124
- “Setting Up Servers Automatically Using Data Saved in a Directory” on page 127

Creating and Saving Setup Data

When you want to work with saved setup data, determine a strategy for naming, encrypting, storing, and serving the data.

One way to create setup data is to use Server Assistant’s offline mode, which lets you work with setup data without connecting to specific servers. You specify setup data, then save it in a file or directory accessible from target servers, as the next two sections describe. Target servers where Mac OS X Server v10.5 software has been installed automatically detect the presence of the saved setup information and use it to set themselves up.

You can define generic setup data that can be used to set up *any* server. For example, you might want to define generic setup data for a server that’s on order, or to configure 50 Xserve computers you want to be identically configured. Alternatively, you can save setup data that’s specifically tailored for a particular server.

Important: When you perform an upgrade installation, make sure that 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, see “How a Server Searches for Saved Setup Data” on page 124.

Setup Data Saved in a File

When you save setup data in a file, a target server detects and uses the file if:

- Setup data the target server recognizes isn't found in a directory the server is configured to use. For information about how a server detects and uses directory data to set itself up, see "Setup Data Saved in a Directory" on page 120.
- The setup file is on a volume mounted locally in `/Volumes/*/Auto Server Setup/`, where `*` is any device mounted under `/Volumes`. A target server searches through volumes alphabetically by device name.

The device that is mounted as a file system can be the server's hard disk or an iPod, DVD, CD, FireWire drive, USB drive, or other device plugged in to the server (for example, `/Volumes/AdminiPod/Auto Server Setup/myserver.example.com.plist`).

- The setup file name is one of the following, when searching for setup files, target servers search for names in the order listed.

`<MAC-address-of-server>.plist` (include leading zeros but omit colons. for example, `0030654dbcef.plist`).

`<IP-address-of-server>.plist` (for example, `10.0.0.4.plist`).

`<partial-DNS-name-of-server>.plist` (for example, `myserver.plist`).

`<built-in-hardware-serial-number-of-server>.plist` (first 8 characters only, for example, `ABCD1234.plist`).

`<fully-qualified-DNS-name-of-server>.plist` (for example, `myserver.example.com.plist`).

`<partial-IP-address-of-server>.plist` (for example, `10.0.plist` matches `10.0.0.4` and `10.0.1.2`).

`generic.plist` (a file that any server will recognize, used to set up servers that need the same setup values). If the serial number specified in the file isn't site licensed, after setup you need to manually set it. Use Server Admin or the following command in the Terminal application: `serversetup -setServerSerialNumber`.

- The correct passphrase is provided to the server when setup data is encrypted. You can use Server Assistant to supply a passphrase interactively, or you can supply the passphrase in a text file. Place the passphrase file on a volume mounted locally on the target server in `/Volumes/*/Auto Server Setup/<pass-phrase-file>`.

The passphrase file can have one of these names. Target servers search for names in the order listed.

`<MAC-address-of-server>.pass` (include leading zeros but omit colons, for example, `0030654dbcef.pass`).

`<IP-address-of-server>.pass` (for example, `10.0.0.4.pass`).

`<partial-DNS-name-of-server>.pass` (for example, `myserver.pass`).

`<built-in-hardware-serial-number-of-server>.pass` (first 8 characters only, for example, `ABCD1234.pass`).

<fully-qualified-DNS-name-of-server>.pass (for example, myserver.example.com.pass).

<partial-IP-address-of-server>.pass (for example, 10.0.pass matches 10.0.0.4 and 10.0.1.2).

generic.pass (a file that any server will recognize). If the server software serial number isn't site licensed, after setup you need to manually set it. Use Server Admin or the following command in Terminal: `serversetup -setServerSerialNumber`.

If you want to reuse saved setup data after reinstalling a server, you can store the server's setup files in a small local partition that isn't erased when you reinstall the server. The setup files are detected and reused after each reinstallation.

Setup Data Saved in a Directory

Using this approach offers the most unattended way to set up multiple servers but it requires that you have a DHCP and directory infrastructure in place.

Using Server Assistant, you save setup data to an existing directory that the computer you're using is configured to access and that you want newly installed servers to retrieve setup data from. The schema of the directory must support stored setup data. Open Directory has built-in support for stored setup data. If you want to store setup data in a non-Apple directory, you first must extend its schema as *Open Directory Administration* describes.

When you save setup data in a directory, a target server detects and uses the setup data if:

- The target server receives its network names (host name, computer name, and local hostname) and its port configuration from a DHCP server.
- The DHCP server is configured to identify the IP address of the directory server where the setup data resides. For DHCP server configuration instructions, see *Network Services Administration*.
- The directory and DHCP servers are running.
- The setup data is stored in the directory in a path named `/AutoServerSetup/` and a record having one of the following names. Target servers search for names in the order listed.

<MAC-address-of-server> (include leading zeros but omit colons, for example, 0030654dbcef).

<IP-address-of-server> (for example, 10.0.0.4).

<partial-DNS-name-of-server> (for example, myserver).

<built-in-hardware-serial-number-of-server> (first 8 characters only, for example, ABCD1234).

<fully-qualified-DNS-name-of-server> (for example, myserver.example.com).

<partial-IP-address-of-server> (for example, 10.0 matches 10.0.0.4 and 10.0.1.2).

generic (a record that any server will recognize, used to set up servers that need the same setup values). If the serial number specified in the file isn't site licensed, after setup you need to manually set it. Use Server Admin or the following command in the Terminal application: `serversetup -setServerSerialNumber`.

- The correct passphrase is provided to the server (setup data stored in a directory should always be encrypted).

You can use Server Assistant to supply a passphrase interactively or you can supply the passphrase in a text file. Place the passphrase file on a volume mounted locally on the target server in `/Volumes/*/Auto Server Setup/<pass-phrase-file>`, where `*` is any device mounted under `/Volumes`. A target server searches through volumes alphabetically by device name.

The passphrase file can have one of the following names. Target servers search for names in the order listed.

<MAC-address-of-server>.pass (include leading zeros but omit colons, for example, 0030654dbcef.pass).

<IP-address-of-server>.pass (for example, 10.0.0.4.pass).

<partial-DNS-name-of-server>.pass. (for example, myserver.pass).

<built-in-hardware-serial-number-of-server>.pass (first 8 characters only, for example, ABCD1234.pass).

<fully-qualified-DNS-name-of-server>.pass (for example, myserver.example.com.pass).

<partial-IP-address-of-server>.pass (for example, 10.0.pass matches 10.0.0.4 and 10.0.1.2).

generic.pass (a file that any server will recognize). If the server software serial number isn't site licensed, after setup you need to manually set it. Use Server Admin or the following command in Terminal: `serversetup -setServerSerialNumber`.

Keeping Backup Copies of Saved Setup Data

Saved setup data isn't only useful for automating the setup of multiple servers. It also provides a way to set up servers again if you need to reinstall server software on them.

You can keep backup copies of setup data files on a network file server. Alternatively, you can store setup data files in a local partition that won't be erased when you reinstall server software.

Using Encryption with Setup Data Files

By default, saved setup data is encrypted for extra security. Before server sets itself up using encrypted data, it must have access to the passphrase used when the data was encrypted.


The passphrase can be provided either interactively (using Server Assistant) or in a file on a local volume of the target server. For example, you can store the file with the passphrase on an iPod, then plug the iPod into each server that needs the passphrase. A server with the IP address 10.0.0.4 would use `/Volumes/MyIPod/Auto Server Setup/10.0.0.4.pass`.

Providing Setup Data Files to Servers


Using Files in the File System

When you place a setup file on a volume (such as a CD, DVD, iPod, USB solid-state drive, or disk partition) mounted locally on a server you've installed but not set up, the server detects the file and uses it to set itself up.

For example, you could store multiple setup files on an iPod and then plug the iPod into the first server that a setup file exists for:



Then you could plug the iPod into the next server.



Each target server recognizes its own file, because it's been named using one of its identifiers and resides in a known location.

For example, a server with WXYZ1234 as the first eight characters of its built-in serial number would use this setup file to set itself up: `/Volumes/MyIPod/Auto Server Setup/WXYZ1234.plist`.

A server's IP address can also be used as an identifier. A server with the IP address of 10.0.0.4 would use the following file: /Volumes/MyIPod/Auto Server Setup/10.0.0.4.plist.

You could also use a single file, which you'd name "generic.plist," to set up multiple servers if the setup data does *not* need to be unique and the servers' network identities are provided using DHCP.


For more information about setup file naming and use, see "How a Server Searches for Saved Setup Data" on page 124.

Using Settings in the Directory

A target server can set itself up using setup data you've stored in a directory the server is configured to access.

Although storing setup data in a directory is the most automated way to set up multiple servers, this approach requires that you set up an infrastructure first so that target servers can locate the setup data stored in the directory.

The most critical components of the infrastructure are DHCP and Open Directory, as the following picture illustrates:



The Open Directory server in this example hosts an LDAP directory where setup data has been saved. The address of the Open Directory server is registered with DHCP service, running on another server in this example. The DHCP service provides the Open Directory server address to the target servers when it assigns IP addresses to those servers. The target servers detect setup data that has been stored for them in the LDAP directory and use it to set themselves up.

You can save setup data in an Apple OpenLDAP directory or in another directory that supports Apple's schema extensions for saved setup data, documented in *Open Directory Administration*.

For more information about setup file naming and use, see "How a Server Searches for Saved Setup Data" on page 124.

How a Server Searches for Saved Setup Data

A freshly installed server sets itself up using saved setup data it finds while using the following search sequence. When the server finds saved setup data that matches the criteria described, it stops searching and uses the data to set itself up.

- 1 The server searches through locally mounted volumes for setup files in `/Volumes/*/Auto Server Setup/`, where `*` is a file system (device) name.

It searches through volumes alphabetically by device name, looking for a file with the extension `.plist` that's named using its MAC address, its IP address, its partial DNS name, its built-in hardware serial number, its fully qualified DNS name, its partial IP address, or `generic.plist`, in that order.

- 2 Next, the server looks in a directory it's configured to use for a setup record in a path named `"AutoServerSetup"`.

It searches for records named using its MAC address, its IP address, its partial DNS name (myserver), its built-in hardware serial number, its fully qualified DNS name (myserver.example.com), its partial IP address, or `"generic,"` in that order.

If the setup data is encrypted, the server needs the correct passphrase before setting itself up. You can use Server Assistant to supply the passphrase interactively, or you can supply the passphrase in a text file in `/Volumes/*/Auto Server Setup/<pass-phrase-file>`.

The target server searches through volumes alphabetically by file system name, looking for a file with the extension `.pass` that's named using its MAC address, its IP address, its partial DNS name, its built-in hardware serial number, its fully qualified DNS name, its partial IP address, or `generic`, in that order.

Important: When you perform an upgrade, make sure that saved setup data won't be inadvertently detected and used by the server you're upgrading. If saved setup data is used, existing server settings are overwritten by the saved settings.

The next two sections provide more details about how to use saved setup data.

Setting Up Servers Automatically Using Data Saved in a File

After you install server software, you can set up the server automatically using data saved in a file.

To save and apply setup data from a file:

- 1 For each server to set up, fill out the "Mac OS X Server Advanced Worksheet" on page 197.
- 2 On an administrator computer, open Server Assistant in `/Applications/Server/`.

You don't need to be an administrator on the administrator computer to use Server Assistant.

- 3 In the Welcome pane, select “Save advanced setup information in a file or directory record” to work in offline mode, which doesn’t require a server connection.
- 4 In the Language pane, specify the language you want to use to administer the target servers.
- 5 If you want to create a setup file, go to step 6; if you want to work with an existing setup file, go to 7.

If you intend to create a generic setup file because you want to use the file to set up more than one server, don’t specify network names (computer name and local hostname) and make sure that each network interface (port) is set to be configured using DHCP or BootP.

- 6 Click Continue and enter the setup data as you move through the Assistant panes, following the onscreen instructions.
- 7 In the Language pane, choose File > Open Configuration File or File > Open Directory Record to load the saved setup data you want to work with. If the saved setup data is encrypted, enter the passphrase when prompted.
Optionally choose View > Jump to Review to review the setup data, then use Go Back as necessary to change it.
- 8 In the Network Interfaces pane, click Add to specify network interfaces.
- 9 After you specify setup data, review the summary displayed by Server Assistant and optionally click Go Back to change data.
- 10 Click Save As, then select Configuration File.
- 11 To encrypt the file, select Save in Encrypted Format, and then enter and verify a passphrase.

You must supply the passphrase before a target server can use an encrypted setup file.

- 12 Click OK, navigate to the location where you want to save the file, name the file using one of the following options, and click Save.

When searching for setup files, target servers search for names in the order listed.

<MAC-address-of-server>.plist (include leading zeros but omit colons, for example, 0030654dbcef.plist)

<IP-address-of-server>.plist (for example, 10.0.0.4.plist)

<partial-DNS-name-of-server>.plist (for example, myserver.plist)

<built-in-hardware-serial-number-of-server>.plist (first 8 characters only, for example, ABCD1234.plist)

<fully-qualified-DNS-name-of-server>.plist (for example, myserver.example.com.plist)

<partial-IP-address-of-server>.plist (for example, 10.0.plist matches 10.0.0.4 and 10.0.1.2).

generic.plist (a file that any server will recognize, used to set up servers that need the same setup values)

- 13 Place the file in a location where target servers can detect it.

A server can detect a setup file if it resides on a volume mounted locally in `/Volumes/*/Auto Server Setup/`, where `*` is a device mounted under `/Volumes`. The device can be the server's hard disk or an iPod, DVD, CD, FireWire drive, USB drive, or other device plugged into the server.

For example, if you have an iPod named AdminiPod, the path used would be `/Volumes/AdminiPod/Auto Server Setup/<setup-file-name>`.

- 14 If the setup data is encrypted, make the passphrase available to target servers.

You can supply the passphrase interactively using Server Assistant, or you can provide it in a text file.

To provide the passphrase in a file, use step 15. To provide it interactively, use step 16.

- 15 To provide a passphrase in a file, create a text file and enter the passphrase for the saved setup file on the first line, and then save the file using one of the following names. Target servers search for names in the order listed.

`<MAC-address-of-server>.pass` (include leading zeros but omit colons, for example, `0030654dbcef.pass`)

`<IP-address-of-server>.pass` (for example, `10.0.0.4.pass`)

`<partial-DNS-name-of-server>.pass`. (for example, `myserver.pass`)

`<built-in-hardware-serial-number-of-server>.pass` (first 8 characters, for example, `ABCD1234.pass`)

`<fully-qualified-DNS-name-of-server>.pass` (for example, `myserver.example.com.pass`)

`<partial-IP-address-of-server>.pass` (for example, `10.0.pass` matches `10.0.0.4` and `10.0.1.2`)

`generic.pass` (a file that any server will recognize)

Save the passphrase file on a volume mounted locally on the target server in `/Volumes/*/Auto Server Setup/<pass-phrase-file>`, where `*` is a device mounted under `/Volumes`.

- 16 To provide a passphrase interactively, use Server Assistant on an administrator computer that can connect with the target server.

- a In the Welcome or Destination pane, choose File > Supply Passphrase.

- b In the dialog box, enter the target server's IP address, password, and passphrase.

- c Click Send.

- 17 If you're using a generic setup file and the serial number isn't site licensed, after setup you must specify the server's serial number by using Server Admin or the command line.

In Server Admin, select the server, click Settings, and click General. Alternatively, in the Terminal application, use `ssh` to connect with the server and enter `serversetup - setServerSerialNumber`.

For a description of the layout of a saved setup file and more information about the `serversetup` command, see *Command-Line Administration*.

Setting Up Servers Automatically Using Data Saved in a Directory

After you install server software, you can set it up automatically using data saved in a directory. This method requires a preexisting directory and DHCP infrastructure, as the procedure below describes.

To save and apply setup data in a directory record:

- 1 Make sure the directory where you want to save setup data exists, that its schema supports stored setup data, and that it's accessible from the administrator computer you're using.

Open Directory Administration describes how to set up and access directories. It also describes the schema for stored setup data. Stored setup data support is built into Apple OpenLDAP directories, but you must extend the schema of other directories to support stored setup data.

- 2 Fill out the "Mac OS X Server Advanced Worksheet" for each server you want to set up. The "Mac OS X Server Advanced Worksheet" on page 197.
- 3 On an administrator computer, open Server Assistant in `/Applications/Server/`. You don't need to be an administrator on the administrator computer to use Server Assistant.
- 4 In the Welcome pane, select "Save advanced setup information in a file or directory record" to work in offline mode, which doesn't require a server connection.
- 5 In the Language pane, specify the language you want to use to administer the target server.
- 6 0.If you want to create a setup, go to step 7; if you want to work with an existing setup, use step 8.

If you're creating generic setup data, don't specify network names (computer name and local hostname) and make sure that each network interface (port) is set to be configured using DHCP or BootP.

- 7 Click Continue and enter the setup data as you move through the Assistant panes, following the onscreen instructions.
- 8 In the Language pane, choose File > Open Configuration File or File > Open Directory Record to load the saved setup data you want to work with.
If the saved setup data is encrypted, enter the passphrase when prompted.
Optionally choose View > Jump to Review to review the setup data, then use Go Back as necessary to change it.
- 9 In the Network Interfaces pane, click Add to specify network interfaces.

10 After you specify setup data, review the summary displayed by Server Assistant and optionally click Go Back to change data.

11 Click Save As, then select Directory Record.

12 To encrypt the file, select Save in Encrypted Format, and then enter and verify a passphrase.

You must supply the passphrase before a target server can use an encrypted directory record.

13 Specify the directory where you want to save the setup in, name the setup record, and click OK; then, when prompted, enter information required to authenticate as a directory domain administrator.

Settings are saved in the directory in AutoServerSetup.

Target servers search for record names in the following order:

<MAC-address-of-server> (include leading zeros but omit colons. for example, 0030654dbcef)

<IP-address-of-server> (for example, 10.0.0.4)

<partial-DNS-name-of-server> (for example, myserver)

<built-in-hardware-serial-number-of-server> (first 8 characters, for example, ABCD1234)

<fully-qualified-DNS-name-of-server> (for example, myserver.example.com)

<partial-IP-address-of-server> (for example, 10.0 matches 10.0.0.4 and 10.0.1.2)

generic (a record that a server will recognize, used to set up servers that need the same setup values).

14 Make sure the infrastructure is in place so the servers you want to use the stored setup record for can find it.

The directory server storing the setup record must be running. DHCP must be configured to identify the directory server to the target servers using Option 95. In addition, you might need to configure DNS if your directory data includes DNS names.

For additional infrastructure information, see "Defining Server Setup Infrastructure Requirements" on page 30. *Open Directory Administration* and *Network Services Administration* provide instructions for setting up directories and DHCP.

15 If the setup data is encrypted, make the passphrase available to the target server.

You can supply the passphrase interactively, using Server Assistant, or you can provide it in a text file.

To provide the passphrase in a file, use step 16; to provide it interactively, use step 17.

16 To provide a passphrase in a file, create a text file and enter the passphrase for the saved setup file on the first line, and then save the file using one of the following names.

Target servers search for names in the order listed.

<MAC-address-of-server>.pass (include leading zeros but omit colons, for example, 0030654dbcef.pass)

<IP-address-of-server>.pass (for example, 10.0.0.4.pass)

<partial-DNS-name-of-server>.pass. (for example, myserver.pass)

<built-in-hardware-serial-number-of-server>.pass (first 8 characters, for example, ABCD1234.pass)

<fully-qualified-DNS-name-of-server>.pass (for example, myserver.example.com.pass)

<partial-IP-address-of-server>.pass (for example, 10.0.pass matches 10.0.0.4 and 10.0.1.2)
generic.pass (a file that a server will recognize).

Put the passphrase file on a volume mounted locally on the target server in `/Volumes/* /Auto Server Setup/<pass-phrase-file>`, where `*` is a device mounted under the directory `/Volumes`.

- 17 To provide a passphrase interactively, use Server Assistant on an administrator computer that can connect with the target server.
 - a In the Welcome or Destination pane, choose File > Supply Passphrase.
 - b In the dialog box, enter the target server's IP address, password, and the passphrase.
 - c Click Send.
- 18 If you're using a generic setup record and the server serial number isn't site licensed, specify the server's serial number by using Server Admin or the command-line after setup.

In Server Admin, select the server, click Settings, and click General. To use the command-line, in the Terminal application use `ssh` to connect with the server and enter `serversetup -setServerSerialNumber`.

For a description of the schema of setup data saved in a directory, see *Open Directory Administration*. For information about `serversetup`, see *Command-Line Administration*.

Determining the Status of Setups

Normally, when setup is complete, the server restarts, and it starts up to the login window. If setup isn't successful, there are several methods by which you're notified.

Using the Destination Pane for Setup Status Information

Server Assistant displays error information in its Destination pane. To access this pane, on the Welcome pane select "Set up a remote server" and click Continue.

If the server isn't listed, click Add to list it. Select the server and review the information displayed.

You can save a list of servers you're interested in monitoring in the Destination pane using File > Save Server List. When you want to monitor the status of those servers, choose File > Load Server List.

Handling Setup Failures

When a server's setup fails, an error log is created as /System/Library/ServerSetup/Configured/POR.err on the target server. The contents of this log can be displayed and the log file deleted on a remote administrator computer.

Double-click the error icon for a server on Server Assistant's Destination pane. If prompted, supply the preset password and click Send.

The log contents are displayed, and you can click Delete to delete the log file. Setup can't be reinitiated until this file has been deleted.

If setup fails because a passphrase file can't be found when using setup data saved in a file or directory record, you can:

- Use Server Assistant to supply a passphrase interactively. On the Destination pane, choose File > Supply Passphrase.
- Supply the passphrase in a text file. Place the passphrase file on a volume mounted locally on the target server in /Volumes/*/Auto Server Setup/<pass-phrase-file>, where * is any device mounted under /Volumes/. A target server searches through volumes alphabetically by device name.

If a remote server setup fails for any other reason, reinstall the server software and repeat initial setup.

If a local server setup fails, restart the computer, rerun Server Assistant, and reinitiate setup, or reinstall the server software.

Handling Setup Warnings

When setup completes but a condition that warrants your attention exists, a warning log is created as /Library/Logs/ServerAssistant.POR.status on the target server. Click the target server's desktop link named ServerAssistant.status to open this file.

Here are some messages you may encounter in the log:

- The server software serial number is invalid. Open Server Admin, select the server in the Servers list, click Settings, and click General. Enter the correct serial number, and click Save.
- Because this server was set up using a generic file or directory record and the serial number isn't site licensed, you must enter the server software serial number using Server Admin. Open Server Admin, select the server in the Servers list, click Settings, and click General. Enter the correct serial number, and click Save.
- The server administrator user defined in the setup data already exists on the server you've upgraded.

Getting Upgrade Installation Status Information

When you perform an upgrade, log files may be placed on the target server. For information about upgrade logs, see the information about upgrading in *Upgrading and Migrating*.

Setting Up Services

After setting up an advanced configuration, you must configure services using Server Admin and add users and groups using Workgroup Manager.

The following sections survey initial setup of individual services and tell you where to find instructions for tailoring services to support your needs.

Adding Services to the Server View

Before you can set up services, you must add the service to the server view in Server Admin. For example, by default, no services can be seen for your server. As you select services to administer, the appropriate configuration panes become accessible in a list underneath your computer name.

The first time you launch Server Admin and connect to a newly installed server, you are prompted to select the services you want to set up and configure on that server. When you select the desired services in the list, those services appear underneath the server hostname in the server list.

Before you can enable or configure a service, it must be added to the administered service list.

To change services to administer:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for each service you want to turn on.

Setting Up Open Directory

Unless your server must be integrated with another vendor's directory system or the directory architecture of a server you're upgrading needs changing immediately, you can begin using the directories you configured during server setup.

Open Directory Administration provides instructions for all aspects of Open Directory domain and authentication setup, including:

- Setting up client computer access to shared directory data
- Replicating LDAP directories and authentication information of Open Directory masters
- Integrating with Active Directory and other non-Apple directories
- Configuring single sign-on

- Using Kerberos and other authentication techniques

Setting Up User Management

Unless you're using a server exclusively to host Internet content (such as web pages) or perform computational clustering, you probably want to set up user accounts in addition to the administrator accounts created during server setup.

User Management tells you how to use Workgroup Manager to connect to the directory, define user settings, set up group accounts and computer lists, define managed preferences, and import accounts.

To set up a user account:

- 1 Open Workgroup Manager.
- 2 Authenticate to the directory as the directory administrator.
- 3 At the top of the application window, click the Accounts button to select the directory you want to add users to.
- 4 Click the New User button.
- 5 Specify user settings in the panes that appear.

You can set up user accounts by using Workgroup Manager to import settings from a file.

Setting Up File Services

When you turn on file sharing services, users can share items in selected folders. You enable and configure File Services and share points using Server Admin. In versions of Mac OS X Server before to Leopard server, share points were created using Workgroup Manager. This functionality has now migrated to Server Admin.

File Services Administration provides instructions for creating, configuring, and managing share points for file sharing using all the protocols.

To set up file sharing:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the appropriate checkbox for each file service you want to turn on.
 - To share with Macintosh computers, turn on Apple file service (AFP service).
 - To share with Windows computers, turn on SMB service.
 - To provide File Transfer Protocol (FTP) access, turn on FTP service.
 - To share with UNIX computers, turn on NFS service.
- 4 Select File Sharing in the toolbar.
- 5 Select a volume or folder you want to share.

- 6 Select “Share this item” for each folder or volume you want to share.
- 7 Click the other tabs to specify attributes for the share point.

Setting Up Print Service

When you turn on print service, server users can share network PostScript printers or Postscript and non-Postscript printers connected directly to the server.

A queue is set up automatically for any USB printer connected to the server. No printer queues are set up automatically for network printers, but they’re easy to add.

To set up a shared printer queue:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for Print service.
- 4 In the list beneath the server, select Print service.
If it isn’t running, click the Start Print button.
- 5 Click Queues.
- 6 Click the Add (+) button.
- 7 Choose a connection protocol, identify a printer, then click OK.

Users of Mac OS X computers can now add the printer using Printer Setup Utility.

For more information about setting up print services, see *Print Service Administration*.

Setting Up Web Service

You can use the Apache HTTP Server that comes with Mac OS X Server to host server and user websites.

If you turned on web service in Server Assistant, your server can begin serving HTML pages from server and user folders.

- To view the main server site, open a web browser on any computer with access to the server and enter the server’s IP address or domain name.
- To view a user site, add a slash (/), a tilde (~), and the user’s short name after the server address. For example, enter

```
http://192.268.2.1/~someuser
```

To turn on web service if it’s not running:

- 1 If you have the HTML files for your main site, copy them into the Documents folder in the /Library/WebServer/ directory.

If the files that make up your site are organized in folders, copy the entire folder structure to the Documents folder.

For a user site, the files go into the Sites folder in the user’s home folder.

Make sure the web content files and folders have the required permissions and ownership. For normal web access, and for WebDAV Read-Only access, the files must be readable by user `www`, and the folders (including all ancestral folders) must be readable and searchable by user `www`. In addition, for WebDAV Read/Write access, the files must be writable by user named “`www`,” and the immediately enclosing folder must be writable by user named “`www`.”

If you don't have your own HTML files yet, you can still turn on web service to see how it works using the default start pages provided with Mac OS X Server.

- 2 Open Server Admin.
- 3 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 4 Select the checkbox for Web service.
- 5 In the list beneath the server, click the button for web service.
- 6 If it isn't running, click the Start Service button in the toolbar.

Web Technologies Administration describes the many features of web service, including how to set up SSL for a site, enable WebMail, and use WebDAV for file sharing.

Setting Up Mail Service

Providing full mail service for your users requires additional configuration beyond what can be described here. *Mail Service Administration* provides instructions for setting up and managing a mail server.

Setting Up Network Services

If you want a server to host any of the following network services, see *Network Services Administration* for setup instructions:

- DHCP service
- DNS
- Firewall service
- Network Address Translation (NAT)
- RADIUS
- VPN
- Network time service

Setting Up System Image and Software Update Services

For details on using NetBoot and NetInstall to simplify the management and installation of client operating systems and other software, see *System Imaging and Software Update Administration*.

It tells you how to create disk images and set up Mac OS X Server so other Macintosh computers can start up from, or install, those images over the network.

The same guide describes how to set up software update service, which lets you customize updates of Apple software on client computers.

To enable NetBoot and NetInstall service for administration:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for NetBoot service.

Setting Up Media Streaming and Broadcasting

For information about how to manage a streaming server that delivers media streams live or on demand to client computers, see *QuickTime Streaming and Broadcasting Administration*.

To enable QuickTime Streaming Service for administration:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for QuickTime Streaming service.

Setting Up Podcast Producer

For information about how to manage a a podcast production server that delivers syndicated media to client computers, see *Podcast Producer Administration*.

To enable Podcast Producer service for administration:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for Podcast Producer service.

Setting Up WebObjects Service

If you want to develop WebObjects applications, see the WebObjects Reference Library, available at developer.apple.com/referencelibrary/WebObjects/. If you want to set up a WebObjects application server, see the Deployment section of the WebObjects Reference Library. More information about WebObject service can be found in *Web Technologies Administration*.

To enable WebObject service for administration:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for WebObject service

Setting Up iChat Service

In addition to services already discussed that help users stay in touch (for example, mail and file services and group accounts and preferences), you can set up an iChat server.

How you use Server Admin to set up iChat service is described in *iChat Service Administration*.

To enable iChat service for administration:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for iChat service.

Setting Up iCal Service

When you turn on iCal Service, you can share and edit calendars for individuals and groups. Using a CalDAV enabled calendar application, you can share, view, and edit calendars with others.

To use Server Admin to set up iCal service, see *iCal Service Administration*.

To enable iCal service for administration:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Services tab.
- 3 Select the checkbox for iCal service.

This chapter shows you how to complete ongoing management for your systems, including setting up administrator computers, designating administrators, and maintaining service uptime.

Sections include:

- “Ports Used for Administration” on page 138
- “Ports Open By Default” on page 138
- “Computers You Can Use to Administer a Server” on page 138
- “Using the Administration Tools” on page 140
- “Opening and Authenticating in Server Admin” on page 140
- “Adding and Removing Servers in Server Admin” on page 141
- “Grouping Servers Manually” on page 142
- “Grouping Servers Using Smart Groups” on page 142
- “Working with Settings for a Specific Server” on page 143
- “Administering Services” on page 147
- “Tiered Administration Permissions” on page 151
- “Workgroup Manager Basics” on page 152
- “Administering Accounts” on page 153
- “Working With Pre-Version 10.5 Computers From Version 10.5 Servers” on page 157
- “Service Configuration Assistants” on page 157
- “Critical Configuration and Data Files” on page 157
- “Improving Service Availability” on page 161
- “Setting Up Your Server for Automatic Reboot” on page 163
- “Load Balancing” on page 170
- “Daemon Overview” on page 171

Ports Used for Administration

For Apple's administration applications to function, the following ports must be enabled.

Port number and type	Tool used
22 TCP	SSH command-line shell
311 TCP	Server Admin (with SSL)
625 TCP	Workgroup Manager
389, 686 TCP	Directory
80 TCP	QuickTime Streaming Management
4111 TCP	Xgrid Admin

In addition, other ports must be enabled for each service you want to run on your server. For a port reference guide, see *Network Services Administration* and the manual for the appropriate service.

Ports Open By Default

After Setup, the firewall is off by default in Advanced Server mode, and therefore all ports are open. When the firewall is turned on, all ports are blocked except the following for all originating IP addresses:

Port number and type	Service
22 TCP	SSH command-line shell
311 TCP	Server Admin (with SSL)
626 UDP	Serial number support
625 TCP	Remote Directory Access
ICMP incoming and outgoing	standard ping
53 UDP	DNS name resolution

Computers You Can Use to Administer a Server

To administer a server locally using the graphical administration applications (in / Applications/Server/) log in to the server as a server administrator and open them.


To administer a remote server, open the applications on an administrator computer. An administrator computer is any Mac OS X Server or Mac OS X v10.5 or later computer where the administration tools have been installed from the *Mac OS X Server Admin Tools* CD. See "Setting Up an Administrator Computer" on page 139.

You can run command-line tools from the Terminal Application (in /Applications/Utilities/) on any Mac OS X Server or Mac OS X computer. You can also run command-line tools from a UNIX workstation.

Setting Up an Administrator Computer

An administrator computer is a computer with Mac OS X or Mac OS X Server v10.5 or later that you use to manage remote servers.

In the following illustration, the arrows originate from administrator computers and point to servers the administrator computers might be used to manage.



When you've installed and set up a Mac OS X Server that has a display, keyboard, and optical drive, it's already an administrator computer. To make a computer with Mac OS X into an administrator computer, you must install additional software.

To enable remote administration of Mac OS X Server from a Mac OS X computer:

- 1 Make sure the Mac OS X computer has Mac OS X v10.5 or later installed, and in addition, make sure the computer has at least 512 MB of RAM and 1 GB of unused disk space.
- 2 Insert the *Mac OS X Server Admin Tools* CD.
- 3 Open the Installer folder.
- 4 Start the installer (ServerAdministrationSoftware.mpkg) and follow the onscreen instructions.

Using a Non-Mac OS X Computer for Administration

You can use a non-Mac OS X computer that offers SSH support, such as a UNIX workstation, to administer Mac OS X Server using command-line tools. For more information, see *Command-Line Administration*.

You can also use any computer which can run a VNC viewer to administer Mac OS X Server. Administering the server via VNC is the same as using the server's keyboard, mouse, and monitor locally.

You enable a VNC server on the Mac OS X Server by enabling Screen Sharing in the Sharing pane of System Preferences.

Using the Administration Tools

Information about administration tools can be found on the pages indicated in the following table.

Use this application or tool	To	See
Installer	Install server software or upgrade it from v10.2 or 10.3.	page 79
Server Assistant	Set up a v10.5 server.	page 112
Workgroup Manager	Administer accounts and their managed preferences.	page 152
Server Admin	Configure and monitor services and administrator access, and configure share points. Set up and manage QuickTime media streaming.	page 143 page 39
System image tools	Manage NetBoot and NetInstall disk images.	page 48
Server Monitor	Monitor Xserve hardware.	page 174
QTSS Publisher	Manage media and prepare it for streaming or progressive download.	page 49
Apple Remote Desktop (optional)	Monitor and control other Macintosh computers.	page 49
Command-line tools	Administer a server using a UNIX command shell.	page 49
Xgrid Admin	Monitor local or remote Xgrid controllers, grids, and jobs.	page 50

You use Server Admin to administer services on one or more Mac OS X Server computers. Server Admin also lets you specify settings that support multiple services, such as creating and managing SSL certificates and specifying which users and groups can access services.

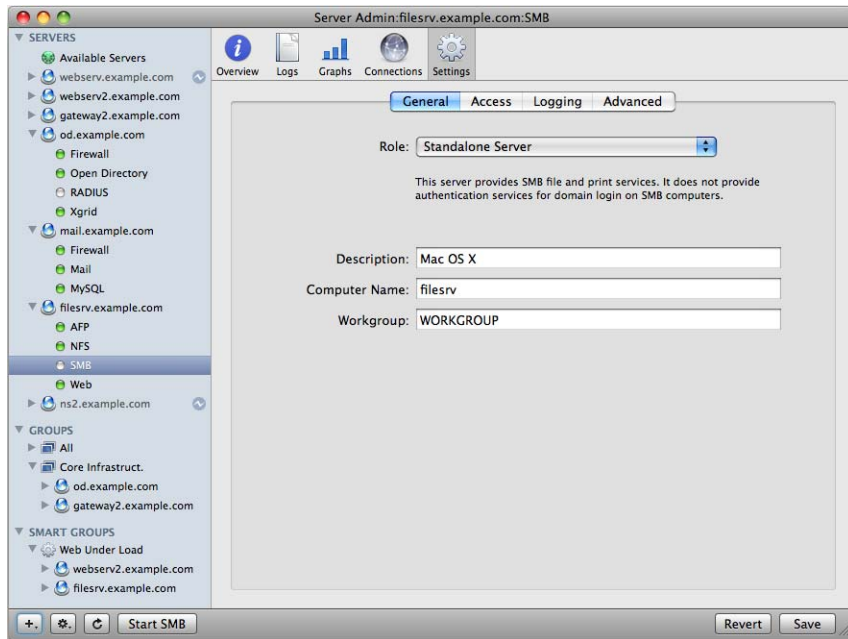
Opening and Authenticating in Server Admin

Server Admin is installed in `/Applications/Server/`. You can open it in the Finder, or you can open it by clicking the Server Admin icon in the Dock, or by clicking the Admin button on the Workgroup Manager toolbar.

To select a server to work with, enter its IP address or DNS name in the login dialog box, or click Browse to choose from a list of servers. Specify the user name and password for a server administrator, then click Connect.

Adding and Removing Servers in Server Admin

The servers you can administer using Server Admin appear in the Servers list on the left side of the application window.



You can add a server to the Servers list and log in to it in two ways:

- Click the Add (+) button in the bottom action bar and choose Add Server.
- Choose Server > Add Server from the menu bar.

The next time you open Server Admin, any server you've added is displayed in the list. To change the order of servers in the list, drag a server to the new location in the list.

You can remove a server from the Servers list in a similar fashion. First you select the server to remove, then you do one of the following:

- Click the Perform Action button in the bottom action bar and choose Disconnect then Remove Server.
- Choose Server > Disconnect, and then choose Server > Remove Server from the menu bar.

If a server in the Servers list appears gray, double-click the server or click the Connect button in the toolbar to log in again. Select the “Remember this password in my keychain” option while you log in to enable auto-reconnect the next time you open Server Admin.

Grouping Servers Manually

Server Admin displays computers in groups in the Server List section of the application’s window. The default server list is called the All Servers list. This is a list of all possible administered computers that you have added and authenticated to. You can create other groups to organize the computers on your network in any way you wish.

Server groups have the following capabilities:

- You can create as many lists as you want.
- Servers can appear in more than one list.
- Groups can be made in any organization scheme you can imagine: geographic, functional, hardware configuration, even color.
- You can click a group name to see a status overview of all servers in the group.

You can make more specific, targeted groups of servers from your All Servers list. First, you can create blank lists and then add servers to them later from the All Servers list.

To create a server group:

- 1 Click the Add (+) button under the Server list at the bottom of the Server Admin window.
- 2 Select Add Group, and name the group.

You can rename groups by clicking the group and letting the mouse hover over the name for a few seconds. The name should become editable.

- 3 Drag the servers from the All Servers group to the newly created group.

Grouping Servers Using Smart Groups

Server Admin displays computers in groups in the Server List section of the application’s window. The default server list is called the All Servers list. This is a list of all possible administered computers that you have added and authenticated to. You can create a server list that automatically populates based on custom criteria. After you create a smart group, any server added to the All Server list (or other specified list) that matches the criteria is dynamically added to the smart group.

You can match any or all of the following criteria:

- Visible services
- Running services

- Network throughput
- CPU utilization
- IP address
- OS version

To create a server smart group:

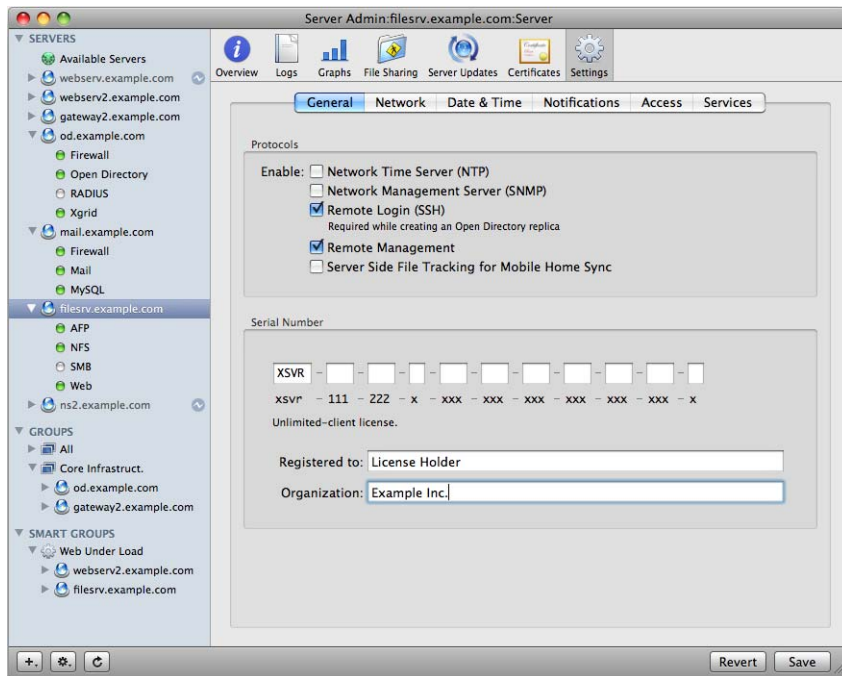
- 1 Click the Add (+) button under the Server list at the bottom of the Server Admin window.
- 2 Select Add Smart Group.
- 3 Name the smart group.
- 4 Define the criteria by which servers will appear in the list, and click OK.

The group will appear in the Server list.

Working with Settings for a Specific Server

To work with general server settings, select a server in the Servers list. You then select from a number of buttons in the toolbar that show configuration options or tabs of configuration options.

The following shows the Settings pane for a server:



The following table contains a summary of what you find for each button:

Toolbar button	Shows
Overview	Information about the server's hardware, software, services, and status
Logs	The system log and security systems log
Graphs	A pictorial history of server activity
Sharing	Configuration options for defining file sharing folders, share points, and automounts
Server Updates	Software updates available from Apple to update the server's software
Certificates	The server's security certificates
Settings	The server's network settings, server software serial number, service access controls, and other information.

When you click Settings, you have access to the following panes:

- **General pane:** Click General to work with the server serial number or to enable Simple Network Management Protocol (SNMP), Network Time Protocol (NTP), Secure Shell (SSH), Remote Management, and server side mobile home sync feature support.

SNMP is a standard that facilitates computer monitoring and management. The server uses the open source net-snmp project for its SNMP implementation.

Although none of the server administration tools use or require SNMP, enabling it enables the server to be monitored and managed from third-party SNMP software such as HP OpenView.

Use the NTP checkbox to enable NTP service. For information about NTP, see *Network Services Administration*.

SSH is a shell you can use to access command-line tools to remotely administer the server with. The Mac OS X Server uses the open source OpenSSH project for its SSH implementation. SSH is also used for other remote server administration tasks, such as initial server setup, Sharing management, and displaying file system paths and the contents of folders in the server administration tools. SSH must be enabled while creating an Open Directory replica, but it can be disabled afterwards.

Remote Management allows the server to be administered by Apple Remote Desktop. You enable and disable Apple Remote Desktop administration in this pane instead of the Sharing pane of System Preferences.

Server side file tracking for mobile home sync is a feature of mobile home folders. For information about when to enable this feature, see *User Management*.

- **Network pane:** Click Network to view or change the server's computer name or local hostname, or to see a list of network interfaces and addressing information for this server.

The computer name is what a user sees when browsing the network (/Network). The local hostname is usually derived from the computer name, but it can be changed.

The network interfaces table shows the name of the interface, the type of addressing (IPv4, or IPv6), the IP address, and the DNS name found by reverse lookup for the address.

- **Date & Time pane:** Click Date & Time to set the server's date and time, NTP source preference, and time zone. For more information about NTP, see *Network Services Administration*.

- **Notifications pane:** Click Notifications to configure Mac OS X Server's automatic event notifications.

You set the mail address and notification trigger in this pane. For more information about notifications, see "Notification in Server Admin" on page 177.

- **Access pane:** Click Access to control user access to some services and to designate administration privileges for users.

When you select the Services tab, you set up access to services to users and groups (referred to as service access control lists, or service ACLs). You can set up the same access to all services, or you can select a service and customize its access settings. Access controls are simple. Choose between enabling all users and groups to use services or enabling only specific users and groups to use services.

When you select the Administrators tab, you designate users to have administration or monitoring privileges for the services on the server. For detailed information about these settings, see "Defining Administrative Permissions" on page 151.

- **Services pane:** Click Services to show or hide services in Server Admin for this server.

Changing the IP Address of a Server

You can change the IP address of a server using the Network pane of System Preferences or the `networksetup` tool.

When a network address change is detected, no matter how the change happened, `changeip` is invoked. The tool `changeip` goes through all configuration files and places where the Server's IP address is stored, and changes the address to conform to the new address. The server's IP address can be changed without `changeip` being invoked from the command-line.

Changing the Server's Host Name After Setup

When you perform an initial server setup for new installations, Server Assistant sets the host name value by assigning AUTOMATIC to the hostname parameter in `/etc/hostname`. This setting causes the server's host name to be the first name that's true in this list:

- The name provided by the DHCP or BootP server for the primary IP address
- The first name returned by a reverse DNS (address-to-name) query for the primary IP address
- The local hostname
- The name "localhost"

After initial setup, if you want to change the host name, don't use the System Preferences Sharing pane to modify the server's computer name; use the `changeip` command-line tool.

For details, see *Command-Line Administration* or the man page for `changeip`.

Changing Server Configuration Type

If you have installed a standard or workgroup configuration server, you can change the server type to an advanced configuration server. All settings you previously set with Server Preferences are retained in the new configuration. No automatic provisioning of user's services occur.

However, you must change the services access controls (SACLs) for services you configured on your standard or workgroup server. For example, if you configured AFP using Server Preferences, you must change the SACLs for AFP using Server Admin to permit access to AFP.

The Server Preferences firewall is separate from the Server Admin firewall, and converting to advanced configuration server disables the Server Preferences firewall. You must enable and configure the firewall accessed through Server Admin.

After conversion, you use Server Admin and the other related tools to administer your server. Server Preferences cannot be used. This is a one-way, one-time conversion.

To change your server configuration:

- 1 Set up an administration computer, which has Server Admin, Workgroup Manager, and other administrative tools installed.

For instructions, see "Setting Up an Administrator Computer" on page 139.

- 2 Launch Server Admin and log in to the switching server.

For instructions on logging in, see "Opening and Authenticating in Server Admin" on page 140.

A dialog sheet appears, asking if you intend to convert the server configuration mode to Advanced.

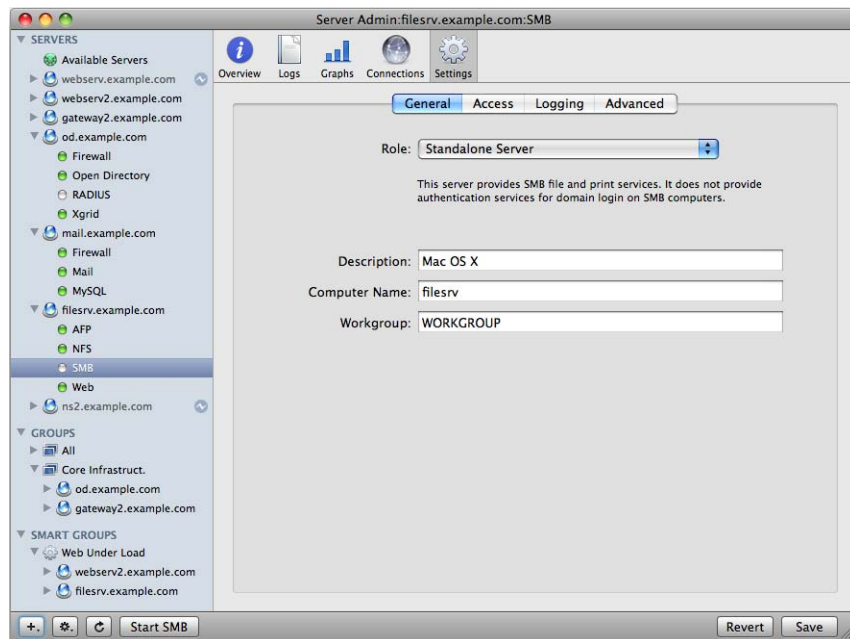
3 Click “Convert to Advanced.”

The server is now no longer in standard or workgroup configuration mode.

Administering Services

To work with a particular service on a server selected in the Servers list of Server Admin, click the service in the list under the server. You can view information about a service (logs, graphs, and so forth) and manage its settings.

The following is a sample service configuration pane in Server Admin.



To start or stop a service, select it and then click Start <service name> or Stop <service name> in the bottom action bar.

Adding and Removing Services in Server Admin

Server Admin can only show you the services you are administering, hiding all other service configuration panes until needed. Before you can administer a service, it must be enabled for the specific server; then that service appears under the server name in the main Server list.

To add or remove a service in Server Admin:

- 1 Select the server that will host the desired service.
- 2 Click the Settings button in the toolbar.
- 3 Click Services.
- 4 Select the desired service, and click Save.

The service now appears in the list, ready for configuration.

Importing and Exporting Service Settings

To copy service settings from one server to another or to save service settings in a property-list file for reuse later, use the Export Service Settings command in Server Admin.

To export settings:

- 1 Select the desired server.
- 2 Choose Server > Export > Service Settings from the menu bar.
- 3 Select the services whose settings you want to copy.
- 4 Click Save.

The file that was created contains all service configuration information as a plist XML document.

To import settings:

- 1 Select the target server to receive the settings.
- 2 Choose Server > Import > Service Settings from the menu bar.
- 3 Find and select the saved service file.

The only file you can use with this function is a properly formatted XML-based plist file, like the one generated from the settings export.

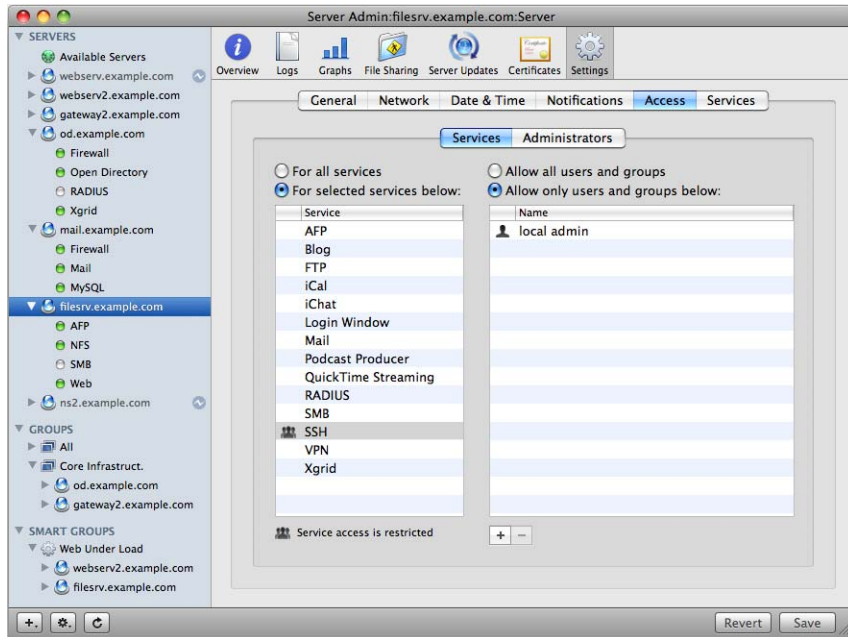
- 4 Click Open.

Controlling Access to Services

You can use Server Admin to configure which users and groups can use services hosted by a server. You set up access to services to users and groups (SACLs). You can set up the same access to all services, or you can select a service and customize its access settings.

Access controls are simple. Choose between allowing all users and groups use services or allowing only selected users and groups use services.

The following shows the Service Access Control List pane in Server Admin:



Select a server in the Servers list, click Settings, click Access, then click Services.

You can separately specify access controls for individual services, or you can define one set of controls that applies for all services that the server hosts.

Using SSL for Remote Server Administration

You can control the level of security of communications between Server Admin and remote servers by choosing Server Admin > Preferences.

By default, Server Admin treats all communications with remote servers as encrypted using SSL. This uses a self-signed 128-bit certificate installed in `/etc/servermgrd/ssl.crt` when you install the server. Communications use HTTPS (port 311). If this option isn't possible, HTTP (port 687) is used and clear text is sent between Server Admin and the remote server.

If you want a greater level of security, also select "Require valid digital signature (SSL)." By default, "Require valid digital signature (SSL)" is disabled. This option uses an SSL certificate installed on a remote server to ensure that the remote server is a valid server.

Before enabling this option, use the instructions in “Requesting a Certificate From a Certificate Authority” for generating a Certificate Signing Request (CSR), obtaining an SSL certificate from an issuing authority, and installing the certificate on each remote server. Instead of placing files in /etc/httpd/, place them in /etc/servermgrd/. You can also generate a self-signed certificate and install it on the remote server.

You can use Server Admin to set up and manage self-signed or -issued SSL certificates used by mail, web, Open Directory, and other services that support them.

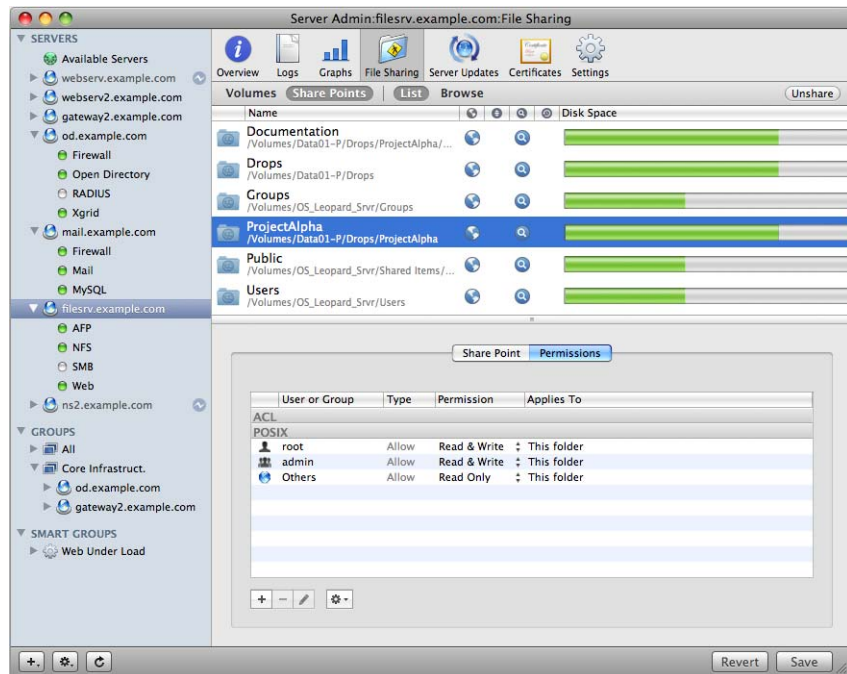
“Certificate Manager in Server Admin” on page 62 provides instructions for using Server Admin to create, organize, and use security certificates for SSL-enabled services. Individual service administration guides describe how to configure specific services to use SSL.

If you’re interested in higher levels of SSL authentication, see the information at www.modssl.org.

Managing Sharing

To work with share points and access control lists, click the File Sharing icon in the Server Admin toolbar. Learn more in *File Services Administration*.

The following is the File Sharing configuration pane in Server Admin.



Tiered Administration Permissions

In previous releases of Mac OS X Server, there were two classes of users: admin and everyone else. Admin users could make any change to the settings of any service or change any directory data as well as passwords and password policies.

In Mac OS X Server v10.5, you can now grant individuals and groups certain administrative permissions, without adding them to the UNIX “admin” group (in other words, you can make them administrator users). There are two levels of permissions:

- **Administer:** This level of permission is analogous to being in the UNIX admin group. You can change any setting on the server for the designated server and service only.
- **Monitor:** This level of permission allows you to view Overview panes, Log panes, and other information panes in Server Admin, as well as general server status data in server status lists. You do not have access to any saved service settings.

Any user or group can be given these permissions for either all services or for only selected services. The permissions are stored on a per-server basis.

The only users that can change the tiered administration access list are users that are truly in the UNIX admin group.

The Server Admin application will update to reflect what operations are possible for a user’s permissions. For example, some services are hidden or the Settings pane is dimmed when you can only monitor that service.

Because the feature is enforced on the server side, the permissions also impact the usage of `serveradmin`, `dscl`, `dsimport`, and `pwdpolicy` command-line tools because all of these tools are limited to the permissions configured for the administrator in use.

Defining Administrative Permissions

You can decide if a user or group can monitor or administer a server or service without giving them the full power of a UNIX administrative user. Assigning effective permissions to users creates a tiered administration, where some but not all administrative duties can be carried out by designated individuals.

To assign permissions:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Access tab.
- 3 Click the Administrators tab.
- 4 Select whether to define administrative permissions for all services on the server or for select services.
- 5 If you choose to define permissions by service, select the appropriate checkbox for each service you want to turn on.

If you define permissions by service, be sure to assign administrators to all the active services on the server.

- 6 Click the Add (+) button to add a user or group from the users and group window.
To remove administrative permissions, select a user or group and click the Remove (-) button.
- 7 For each user or group, select the permissions level next to the user or group name.
You can choose Monitor or Administer.

The capabilities of Server Admin to administer the server are limited by this setting, when the server is added to the Server list.

Workgroup Manager Basics

You use Workgroup Manager to administer the following accounts: user accounts, group accounts, and computer lists. You also use it to set preferences for Mac OS X user accounts, group accounts, computers, and access the Inspector, an advanced feature that lets you do raw editing of Open Directory entries.

The following topics describe general Workgroup Manager usage. Instructions for conducting specific administration tasks are available in Workgroup Manager help and in several guides:

- *User Management* tells you how to use Workgroup Manager for managing user accounts, group accounts, computer lists, preferences, and how to import and export accounts.
- *File Services Administration* explains how to use Sharing in Workgroup Manager to manage share points.
- *Open Directory Administration* provides information about using the Inspector.

Opening and Authenticating in Workgroup Manager

Workgroup Manager is installed in /Applications/Server/, you can open it in the Finder, the Dock, or you can open Workgroup Manager by selecting View > Workgroup Manager in the menu bar of Server Admin:

- When you open Workgroup Manager on the server you're using without authenticating, you have read-only access to information displayed in the local domain. To make changes, click the lock icon to authenticate as a server administrator.

This approach is most useful when you're administering various servers and working with several directory domains.

- To authenticate as an administrator for a server, local or remote, enter the server's IP address or DNS name in the login dialog box, or click the directory path area of the Workgroup Manager window to choose another directory server. Specify the user name and password for an administrator of the server, then click Connect.

Use this approach when you'll be working most of the time with a particular server.

After opening Workgroup Manager, you can open a Workgroup Manager window for a different computer by clicking New Window in the toolbar or choosing Server > Connect.

Important: When you connect to a server in Workgroup Manager, make sure the long or short user name you specify matches the capitalization in the user account.

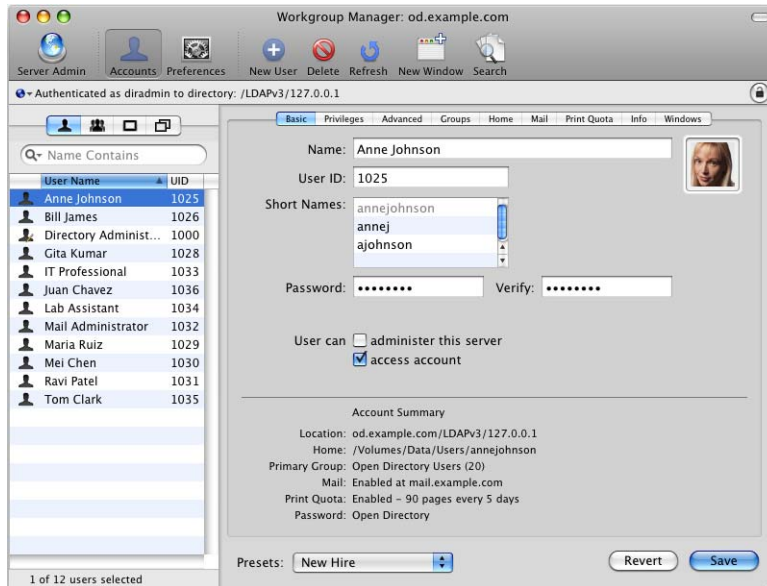
Administering Accounts

User accounts and group memberships are not administered in Server Admin. You need to use Workgroup Manager to add and remove users and groups. For information about account administration, see *User Management*. What follows is a brief synopsis of account administration using Workgroup Manager. Do not use this section as your only source of information about accounts.

Working with Users and Groups

After you log in to Workgroup Manager, the account window appears, showing a list of user accounts. Initially, accounts listed are those stored in the last directory node of the server's search path. When you use other Workgroup Manager windows, such as Preferences, click Accounts in the toolbar to return to the account window.

The following is a sample user record configuration pane in Workgroup Manager:



To specify the directories that store accounts you want to work with, click the small globe icon. To work with different accounts in different Workgroup Manager windows, click New Window in the toolbar.

To administer the accounts listed, click the Users, Groups, or Computers, or Computer Groups button on the left side of the window. You can filter the accounts listed by using the pop-up search list above the accounts list. To refresh the accounts list, click the Refresh button in the toolbar.

To simplify defining an account's initial attributes when you create the account, use presets. A preset is an account template.

To create a preset, select an account, set up all the values the way you want them, then choose Save Preset from the Presets pop-up menu at the bottom of the window.

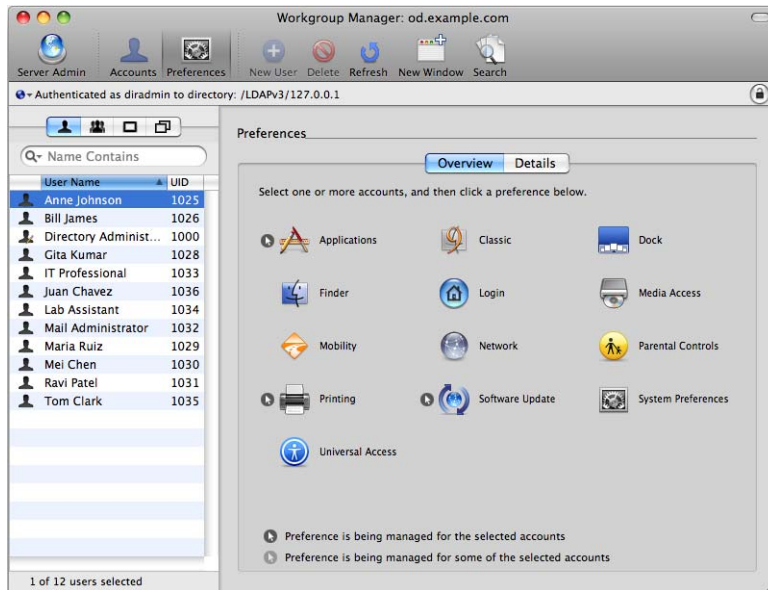
To work with only accounts that meet specific criteria, click Search in the toolbar. The Search features include the option for batch editing selected accounts.

To import or export accounts, select the accounts, then choose Server > Import or Server > Export, respectively.


Defining Managed Preferences

To work with managed preferences for user accounts, group accounts, or computer lists, click the Preferences icon in the Workgroup Manager toolbar.

The following is the User Preference Management Overview pane in Workgroup Manager:



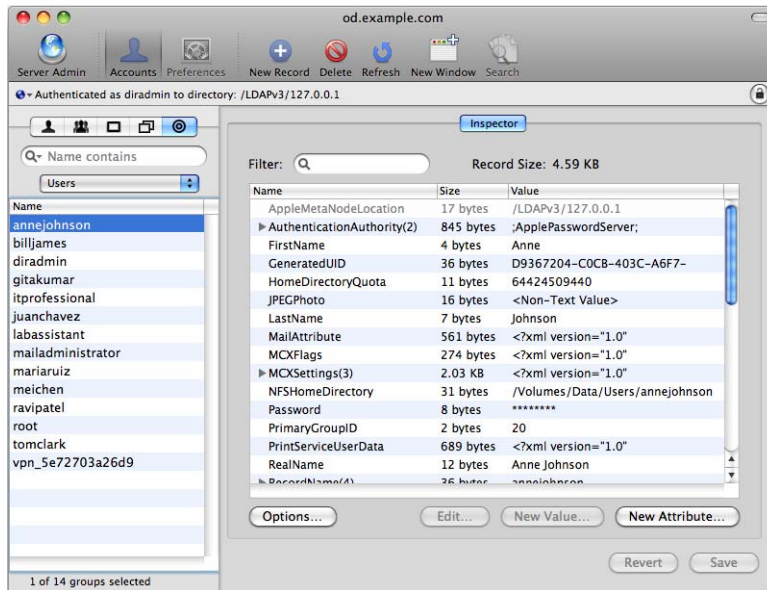
Click Details to use the preference editor to work with preference manifests. The following is a sample of the preference editor sheet in Workgroup Manager:



Working with Directory Data

To work with raw directory data, use Workgroup Manager's Inspector.

The following is the record Inspector pane in Workgroup Manager:



To display the inspector:

- 1 Choose Workgroup Manager > Preferences.
- 2 Enable "Show "All Records" tab and inspector" and click OK.
- 3 Select the "All records" button (which looks like a bull's-eye) to access the Inspector.
- 4 Use the pop-up menu above the Name list to select the records of interest.

For example, you can work with users, groups, computers, share points, and many other directory objects.

Customizing the Workgroup Manager Environment

There are several ways to tailor the Workgroup Manager environment:

- You can control the way Workgroup Manager lists accounts and other behaviors by choosing Workgroup Manager > Preferences.
- To customize the toolbar, choose View > Customize Toolbar.
- To include predefined users and groups in the user and group lists, choose View > Show System Users and Groups.
- To open Server Admin so you can monitor and work with services on particular servers, click the Server Admin icon in the toolbar.

Working With Pre-Version 10.5 Computers From Version 10.5 Servers

You can use the version of Server Admin included with Mac OS X Server v10.5 to administer Mac OS X Server v10.4.11 or later. Workgroup Manager on a v10.5 server can be used to manage Mac OS X clients running Mac OS X v10.3 or later.

After you edit a user record using Workgroup Manager on v10.5, you can only access it using Workgroup Manager on v10.5.

Service Configuration Assistants

Server Admin has configuration assistants to guide you through setting up services that require more setup than a single configuration pane. The assistants present you with all configuration panes necessary to fully enable a service.

Assistants are available for the following services:

- **Gateway Setup:** This assistant helps you set up your server as a network gateway. Launch the assistant using a button in the lower right side of NAT service's Overview page.
- **Mail:** This assistant helps you set up both incoming and outgoing email service. Launch the assistant using a button in the lower right side of Mail service's Overview page.
- **RADIUS:** This assistant helps you set up RADIUS authentication for Apple Airport wireless access points. Launch the assistant using a button in the lower right side of RADIUS service's Overview page.
- **Xgrid:** This assistant helps you set up Xgrid controllers. Launch the assistant using a button in the lower right side of Xgrid service's Overview page.

Critical Configuration and Data Files

When backing up system settings and data, take special care to make sure all your critical configuration files are backed up. The nature and frequency of your backups depend on your organization's backup, archive and restore policies. For more information about creating a backup and restore policy, see "Defining Backup and Restore Policies" on page 32.

The following is a list of configuration and data files for services available on Mac OS X Server.

General

File type	Location
Service states	/System/Library/LaunchDaemons/*
SSH configuration files and host's public / private keys	/etc/ssh/*
System keychain	/Library/Keychains/System.keychain

iCal Service

File type	Location
Configuration files	/etc/calDavd/calDavd.plist
Data	/Library/CalendarServer/Documents/

iChat Server

File type	Location
Configuration files	/etc/jabberd/*
Data	mysqldump jabberd2 > jabberd2.backup.sql

Notifications

File type	Location
Configuration files	/etc/emond.d/ /etc/emond.d/rules/ /Library/Keychains/System.keychain

QuickTime Streaming Server

File type	Location
Configuration files	/Library/QuickTimeStreamingServer/Config/* /Library/QuickTimeStreamingServer/Playlists/* /Library/Application Support/Apple/QTSS Publisher/*
Data: (default locations)	/Library/QuickTimeStreamingServer/Movies/* ~user/Sites/Streaming/*

Firewall Service

File type	Location
Configuration files	/etc/ipfilter/*

NAT Service

File type	Location
Configuration files	/etc/nat/*

Mail Services

The following are the configuration files and data stores for mail services.

Mail—SMTP Server Postfix

File type	Location
Configuration files	/etc/postfix/
Data: (default locations)	/var/spool/postfix/

Mail—POP/IMAP Server Cyrus

File type	Location
Configuration files	/etc/imapd.conf /etc/cyrus.conf
Data: (mail database default location)	/var/imap
(mail data store)	/var/spool/imap

Custom locations are defined in /etc/imapd.conf using the following keys with default values:

Custom locations	Key: Value pair
Mail database location	configdirectory: /var/imap
Mail data store location	partition-default: /var/spool/imap
Additional data store partitions (no default value)	partition-xxx: /var/spool/mail_xxx There can be multiple additional data store partitions

Mail—Amavisd

File type	Location
Configuration files	/etc/amavisd.conf
Data: (default locations)	/var/amavis/

Mail—Clam AV

File type	Location
Configuration files	/etc/clamav.conf /etc/freshclam.conf

File type	Location
Data: (default locations)	/var/clamav/ /var/virusmails/

Mail—Mailman

File type	Location
Configuration files	/var/mailman/
Data: (default locations)	/var/mailman/

Mail—SpamAssassin

File type	Location
Configuration files	/etc/mail/spamassassin/local.cf
Data: (default locations)	/etc/mail/spamassassin/

MySQL Service

File type	Location
Configuration files	There is no config file for MySQL, but the administrator can create one, which should be backed up if present: /etc/my.cnf
Data: (default locations)	/var/mysql/ mysqldump --all-databases > all.sql

PHP

File type	Location
Configuration files	There is no config file for PHP, but the administrator can create one (copying /etc/php.ini.default to /etc/php.ini and modifying it), which should be backed up if present: /etc/php.ini
Data: (default locations)	as designated by administrator

Web Service

File type	Location
Configuration files	/etc/httpd/* (for Apache 1.3) /etc/apache2/* (for Apache 2.2) /etc/webperfcache/* /Library/Keychains/System.keychain
Data: (default locations)	/Library/WebServer/Documents/

File type	Location
	/Library/Logs/WebServer/*
	/Library/Logs/Migration/webconfigmigrator.log (Apache config migration log)

The default location for web content is configurable and is most likely modified and extended to include multiple virtual host content and WebDAV directories.

Note: Log files for web service are a critical source of revenue for some sites and should be considered for backup. The location is configurable and can be determined using Server Admin.

Wiki and Blog Server

File type	Location
Configuration files	/etc/wikid/*
	/Library/Application Support/Apple/WikiServer (wiki themes and template files)
Data: (default locations)	/Library/Collaboration/
Log files: (default location)	/Library/Logs/wikid/*

Improving Service Availability

Eliminating single points of failure and using Xserve and hardware RAID are some of the things that can boost your server availability. Other things you can do range from simple solutions like using power backup, automatic reboot, and ensuring proper operational conditions (for example, adequate temperature and humidity levels) to more advanced solutions involving link aggregation, load balancing, Open Directory replication, and data backup.

Eliminating Single Points of Failure

To improve the availability of your server, reduce or eliminate single points of failure. A single point of failure is any component in your server environment that, if it fails, causes your server to fail.

Some single points of failure include:

- Computer system
- Hard disk
- Power supply

Although it is almost impossible to eliminate all single points of failure, you should minimize them as much as possible. For example, using a backup system and the IP failover in Mac OS X Server eliminates the computer as a single point of failure. Although both the master and backup computers can fail at once or one after the other, the possibility of such an event happening is negligible.

Another way to prevent a computer from failing is to use a backup power source and take advantage of hardware RAID to mirror the hard disk. With hardware RAID, if the main disk fails, the system can still access the same data on the mirror drive, as is the case with Xserve.

Using Xserve for High Availability

Xserve is designed for extra reliability and hence, high availability.

Although you can use desktop systems like the Power Mac G5 or Mac Pro to provide Mac OS X Server services very reliably, Xserve has the following additional features that make it ideal for high availability situations.

- Xserve has eight fans. In the case of a single fan failure, the other fans speed up to compensate, allowing your server to keep running.
- An independent drive architecture isolates the drives electrically, preventing a single drive failure from causing unavailability or performance degradation of the surviving drives—a common problem with multidrive SCSI implementations.
- Xserve uses Error Correction Code (ECC) logic to protect the system from corrupt data and transmission errors.

Each DIMM has an extra memory module that stores checksum data for every transaction. The system controller uses this ECC data to identify single-bit errors and corrects them on the fly, preventing unplanned system shutdowns.

In the rare event of multiple-bit errors, the system controller detects the error and triggers a system notification to prevent bad data from corrupting further operations.

You can set the Server Monitor software to alert you if error rates exceed the defined threshold.

- Xserve has built-in hardware RAID mirroring, which protects your server from failing if the main drive fails.

For more information about Xserve, visit www.apple.com/xserve/.

Using Backup Power


In the architecture of a server solution, power is a single point of failure. If power goes out, your servers go down without warning. To prevent a sudden disruption in services, consider adding a backup source of power.

Depending on your application, you might choose to use a standby electrical generator or Uninterruptible Power Supply (UPS) devices to gain enough time to notify users of an impending shutdown of services.

Using UPS with Xserve

Xserve does not provide serial port connectivity to UPS, but it can monitor UPS power through the network if the UPS unit has a management network card. For more information, check with UPS vendors.

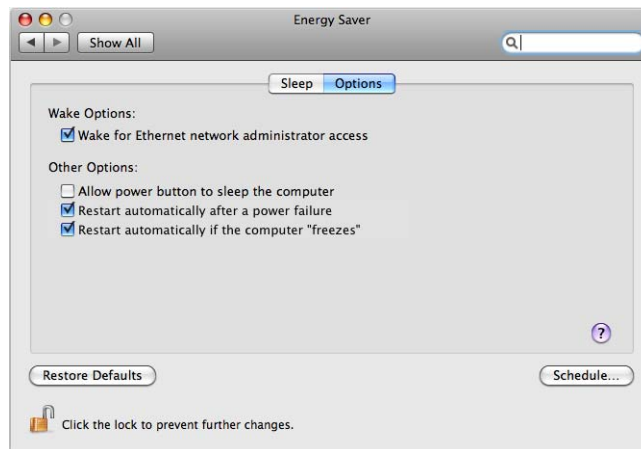
The following illustration is an example of an Xserve connected to a UPS via a network:



Setting Up Your Server for Automatic Reboot

You can set up Energy Saver options on your Mac OS X Server computer to automatically restart if it goes down due to a power failure or system freeze.

The following is the Energy Saver panel of System Preferences:



The automatic reboot options are:

- **Restart automatically after a power failure.** The power management unit automatically starts up the server after a power failure.
- **Restart automatically if the computer freezes.** The power management unit automatically starts up the server after the server stops responding, has a kernel panic, or freezes.

When you select the option to restart after a freeze, Mac OS X Server spawns the `wdticklerd` daemon, which every 30 seconds commands your computer to reboot after 5 minutes. Each time the command is sent, the restart timer is reset. Thus, the timer won't reach 5 minutes as long as the server is running. If the computer does freeze, the power management unit will restart it after 5 minutes.

To enable automatic reboot:

- 1 Log in to the server as an administrator.
- 2 Open System Preferences and click Energy Saver.
- 3 Click Options.
- 4 Under Other Options, select restart options.
- 5 Close System Preferences.

Ensuring Proper Operational Conditions

One factor that can cause your servers to malfunction is overheating. This is especially a problem when you cluster computers in a small space. Other factors such as humidity and power surges can also adversely impact your server.

To protect your servers, make sure you house them in a place where you can control these factors and provide ideal operating conditions. Check the electrical and environmental requirements for your systems to find what these conditions are.

In addition, make sure the facility you deploy your server has a fire alarm, and prepare a contingency plan to deal with this risk.

Providing Open Directory Replication

If you plan to provide Open Directory services, consider creating replicas of your Open Directory master. If the master server fails, client computers can access the replica.

For more information, see the section on setting up Open Directory replicas in *Open Directory Administration*.

Link Aggregation

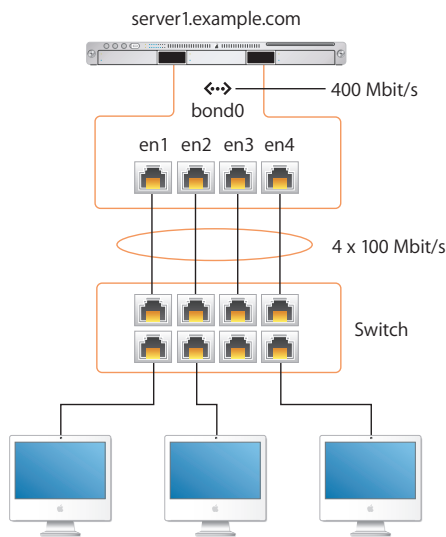
Although not common, the failure of a switch, cable, or network interface card can cause your server to become unavailable. To eliminate these single points of failure, you can use link aggregation or trunking. This technology, also known as IEEE 802.3ad, is built into Mac OS X and Mac OS X Server.

Link aggregation allows you to aggregate or combine multiple physical links connecting your Mac to a link aggregation device (a switch or another Mac) into a single logical link. The result is a fault-tolerant link with a bandwidth equal to the sum of the bandwidths of the physical links.

For example, you can set up an Xserve with four 1-Gbit/s ports (en1, en2, en3, and en4) and use the Network pane of System Preferences to create a link aggregate port configuration (bond0) that combines en1, en2, en3, and en4 into one logical link.

The resulting logical link will have a bandwidth of 4 Gbit/s. This link will also provide fault tolerance. If one or more physical links fail, your Xserve's bandwidth will shrink, but the Xserve can still service requests as long as not all physical links fail at once.

The following illustration shows four Ethernet ports aggregated as a single interface:



Link aggregation also allows you to take advantage of existing or inexpensive hardware to increase the bandwidth of your server. For example, you can form a link aggregate from a combination of multiple 100-Mbit/s links or 1-Gbit/s links.

The Link Aggregation Control Protocol (LACP)

IEEE 802.3ad Link Aggregation defines a protocol called Link Aggregation Control Protocol (LACP) that is used by Mac OS X Server to aggregate (combine) multiple ports into a link aggregate (a virtual port) that can be used for TCP and UDP connections.

When you define a link aggregate, the nodes on each side of the aggregate (for example, a computer and a switch) use LACP over each physical link to:

- Determine whether the link can be aggregated
- Maintain and monitor the aggregation

If a node doesn't receive LACP packets from its peer (the other node in the aggregate) regularly, it assumes that the peer is no longer active and removes the port from the aggregate.

In addition to LACP, Mac OS X Server uses a frame distribution algorithm to map a conversation to a particular port. This algorithm sends packets to the system on the other end of the aggregate only if it has packet reception enabled. In other words, the algorithm won't send packets if the other system isn't listening.

Mapping a conversation to a particular port guarantees that packet reordering will not occur.

Link Aggregation Scenarios


Following are three common aggregation scenarios that you can set up:

- Computer-to-computer
- Computer-to-switch
- Computer-to-switch-pair

These scenarios are described in the following sections.

Computer-to-Computer


In this scenario, you connect the two servers directly (as shown in the following illustration) using the physical links of the link aggregate.



This allows the two servers to communicate at a higher speed without the need for a switch. This configuration is ideal for ensuring back-end redundancy.

Computer-to-Switch

In this scenario shown in the following illustration, you connect your server to a switch configured for 802.3ad link aggregation.




The switch should have a bandwidth for handling incoming traffic equal to or greater than that of the link aggregate (logical link) you define on your server.

For example, if you create an aggregate of four 1-Gbit/s links, you should use a switch that can handle incoming traffic (from clients) at 4 Gbit/s or more. Otherwise, the increased bandwidth advantage in the link aggregate won't be fully realized.

Note: For information about how to configure your switch for 802.3ad link aggregation, see the documentation provided by the switch manufacturer.

Computer-to-Switch-Pair

In this scenario shown in the following illustration, you improve on the computer-to-switch scenario by using two switches to eliminate the switch as a single point to failure:



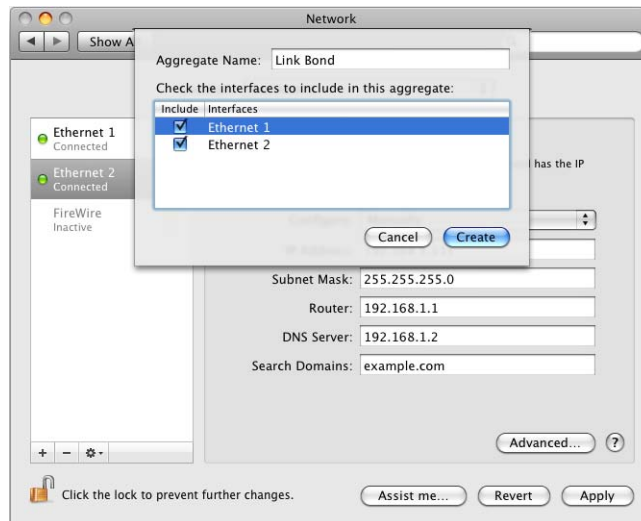
For example, you can connect two links of the link aggregate to the master switch and the remaining links to the backup switch. As long as the master switch is active, the backup switch remains inactive. If the master switch fails, the backup switch takes over transparently to the user.

Although this scenario adds redundancy that protects the server from becoming unavailable if the switch fails, it results in decreased bandwidth.

Setting Up Link Aggregation in Mac OS X Server

To set up your Mac OS X Server for link aggregation, you need a Mac with two or more IEEE 802.3ad-compliant Ethernet ports. In addition, you need at least one IEEE 802.3ad-compliant switch or another Mac OS X Server computer with the same number of ports.

You create a link aggregate on your computer in the Network pane of System Preferences (as shown in the following example):



To create a link aggregate:

- 1 Log in to the server as an administrative user.
- 2 Open System Preferences.
- 3 Click Network.
- 4 Click the Gear button and choose Manage Virtual Interfaces in the pop-up menu.
- 5 Click the Add (+) button, and select New Link Aggregate in the pop-up menu.

Note: You only see this option if you have two or more Ethernet interfaces on your system.

- 6 In the Name field, enter the name of the link aggregate.

- 7 Select the ports to aggregate from the list.
- 8 Click Create.
- 9 Click Done.

By default the system gives the link aggregate the interface name `bond<num>`, where `<num>` is a number indicating precedence. For example, the first link aggregate is named `bond0`, the second is `bond1`, and the third is `bond2`.

The interface name `bond<num>` assigned by the system is different from the name you give to the link aggregate port configuration. The interface name is for use at the command line, but the port configuration name is for use in the Network pane of System Preferences.

For example, if you enter the command `ifconfig -a`, the output refers to the link aggregate using the interface name and *not* the port configuration name:

```
...
bond0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    inet6 fe80::2e0:edff:fe08:3ea6 prefixlen 64 scopeid 0xc
    inet 10.0.0.12 netmask 0xfffff00 broadcast 10.0.0.255
    ether 00:e0:ed:08:3e:a6
    media: autoselect (100baseTX <full-duplex>) status: active
    supported media: autoselect
    bond interfaces: en1 en2 en3 en4
```

You do not delete or remove a link bond from the Network Pane of System Preferences. You remove the bond through the Manage Virtual Interfaces sheet used to create the bond.

Monitoring Link Aggregation Status

You can monitor the status of a link aggregate in Mac OS X and Mac OS X Server using the Status pane of the Network pane of System Preferences.

To monitor the status of a link aggregate:

- 1 Open System Preferences.
- 2 Click Network.
- 3 From the list of network interfaces on the left, choose the link aggregate port virtual interface.
- 4 Click Advanced in the lower right side of the window.
- 5 Select the Bond Status tab.

The Status pane displays a list containing a row for each physical link in the link aggregate. For each link, you can view the name of the network interface, its speed, its duplex setting, the status indicators for incoming and outgoing traffic, and an overall assessment of the status.

Note: The Sending and Receiving status indicators are color-coded. Green means the link is active (turned on) and connected. Yellow means the link is active but not connected. Red means the link can't send or receive traffic.


- 6 To view more information about a link, click the corresponding entry in the list.

Load Balancing

One factor that can cause services to become unavailable is server overload. A server has limited resources and can service a limited number of requests simultaneously. If the server gets overloaded, it slows down and can eventually crash.

One way to overcome this problem is to distribute the load among a group of servers (a server farm) using a third-party load-balancing device. Clients send requests to the device, which then forwards the request to the first available server based on a predefined algorithm. The clients see only a single virtual address, that of the load-balancing device.

Many load-balancing devices also function as switches (as shown in the following illustration), providing two functions in one, which reduces the amount of hardware you need to use.



Note: A load-balancing device must be able to handle the aggregate (combined) traffic of the servers connected to it. Otherwise, the device becomes a bottleneck, which reduces the availability of your servers.

Load balancing provides several advantages:

- **High availability.** Distributing the load among multiple servers helps you reduce the chances that a server will fail due to server overload.
- **Fault tolerance.** If a server fails, traffic is transparently redirected to other servers. There might be a brief disruption of service if, for example, a server fails while a user is downloading a file from shared storage, but the user can reconnect and restart the file download process.
- **Scalability.** If demand for your services increases, you can transparently add more servers to your farm to keep up with the demand.
- **Better performance.** By sending requests to the least-busy servers, you can respond faster to user requests.

Daemon Overview

By the time a user logs in to a Mac OS X system, a number of processes are already running. Many of these processes are known as daemons. A daemon is a background process that provides a service to users of the system. For example, the `cupsd` daemon coordinates printing requests, and the `httpd` daemon responds to requests for web pages.

Viewing Running Daemons

If you want to see the daemons running on your system, use the Activity Monitor application (in `/Applications/Utilities/`). This application lets you view information about all processes, including their resource usage.

You will see the following daemons, regardless of what services are enabled:

- `launchd` (timed job and watchdog process)
- `servermgrd` (administration tool interface process)
- `serialnumberd` (license compliance process)
- `mDNSResponder` (local network service discovery process)

Daemon Control

Although some UNIX-like systems use other tools, Mac OS X Server uses a daemon called `launchd` to control process initialization and timed jobs.

`launchd`

The `launchd` daemon is an alternative to the following common UNIX tools: `init`, `rc`, the `init.d` and `rc.d` scripts, `SystemStarter`, `inetd` and `xinetd`, `atd`, `crond` and `watchdogd`. All of these services should be considered deprecated and administrators are strongly encouraged to move process management duties to `launchd`.

There are two utilities in the `launchd` system: `launchd` daemon and `launchctl` utility.

The launchd daemon also has replaced init as the first process spawned in Mac OS X and is therefore responsible for starting the system at startup. The launchd daemon manages the daemons at both a system and user level. It can:

- Start daemons on demand
- Monitor daemons to make sure they keep running

Configuration files are used by launchd to define the parameters of services and daemons run. The configuration files are property list files stored in the LaunchAgents and LaunchDaemons subdirectories of the Library folders.

For more information about creating the launchd configuration files, see the following Developer Documentation page:

developer.apple.com/documentation/MacOSX/Conceptual/BPSystemStartup/Articles/LaunchOnDemandDaemons.html

The launchctl utility is the command-line tool used to:

- Load and unload daemons
- Start and stop launchd controlled jobs
- Get system utilization statistics for launchd and its child processes
- Set environment settings

Effective monitoring allows you to detect potential problems before they occur and gives you early warning when they occur.

Detecting potential problems allows you to take steps to resolve them before they impact the availability of your servers. In addition, getting early warning when a problem occurs allows you to take corrective action quickly and minimize disruption to your services.

This chapter briefly describes planning a monitoring policy, how to use monitoring tools, and how to find more information.

Planning a Monitoring Policy

Gathering data about your systems is a basic function of good administration. Different types of data gathering are used for different purposes.

- **Historical data collection:** Historical data is gathered for analysis. This could be used for IT planning, budgeting, and getting a baseline for normal server conditions and operations. What kinds of data do you need for these purposes? How long does it need to be kept? How often does it need to be updated? How far in the past does it need to be collected?
- **Real-time monitoring:** Real-time monitoring is for alerts and detecting problems as they happen. What are you monitoring? How often? Does that data tell you what you need to know? Are some of these real-time collections actually for historical purposes?

Planning Monitoring Response

The response to your monitoring is as important as the data collection. In the same way a backup policy is pointless without a restore strategy, a monitoring policy makes little sense without a response policy.

Several factors can be considered for a monitoring response:

- What are appropriate response methods? In other words, how will the response take place?
- What is the time to response? What is an acceptable interval between failure and response?
- What are the scaling considerations? Can the response plan work with all expected (and even unexpected) frequencies of failure?
- Are there testing monitoring systems in place? How do you know the monitoring policy is catching the data you need, and how do you know the responses are timely and appropriate? Have you tested the monitoring system recently?

Server Status Widget

The Server Status Dashboard widget is provided for quick access and information about a single system. The Server Status widget lets you monitor Mac OS X Server v10.5 activity from any computer with Leopard or Leopard Server. Server Status shows you graphs of processor activity, network load, disk usage, polled hourly, daily, or weekly.

You can also see up to six running services and their status reports. By clicking on the service, you can open Server Admin to the appropriate service overview panel.

To configure the Server Status widget:

- 1 Add the widget to the Dashboard like any other widget.
- 2 Enter the server IP address or domain name.
- 3 Supply an administrative or monitoring login name and password.
- 4 Click Done.

To change the server address, login name, or password, click the information button (i) at the top of the widget and change the settings.

Server Monitor

The Server Monitor application can issue alerts via mail, cell phone, or pager notification as soon as it detects critical problems. Built-in sensors detect and report essential operating factors like power, temperature, and the condition of several key components.

The Server Monitor interface allows you to quickly detect problems. In the main window, Server Monitor lists each server on a separate line, with temperature information and the status of each of its components, including fans, disk drives, memory modules, power supplies, and Ethernet connections.

A green status indicator shows the component is OK, a yellow status indicator notes a warning, and a red status indicator notes an error.

Server Monitor works for Xservices only. For more information about Server Monitor, choose Server Monitor Help from Server Monitor's Help menu.

RAID Admin

Like Server Monitor, you can configure RAID Admin to send an email or page when a component is in trouble. For every unit, RAID Admin displays the status of the unit and each of its components, including disk drives, fibre channel, and network connections.

RAID Admin uses green, yellow, or red status indicators. You can also configure it to send you an email or page when a component is in trouble.

In addition, RAID Admin provides you with an overview of the status of the Xserve RAID units that appear in the main window.

For more information about RAID Admin, choose RAID Admin Help from RAID Admin's Help menu.

Console

Use Console to monitor relevant log files for potential problems that might cause your server to fail.

For example, you can monitor your web server's `/var/log/httpd/access_log` file for signs of denial of service attacks. If you detect these signs, you can immediately implement a planned response to prevent your web server from becoming unavailable.

To improve your log monitoring efficiency, consider automating the monitoring process using AppleScript or Terminal commands like `grep` and `cron`. For more information about using `grep` and `cron`, see *Command-Line Administration*

Disk Monitoring Tools

Running out of disk space can cause your server to become unreliable and probably fail. To prevent this from happening, you must constantly monitor disk space usage on your servers and delete or back up files to clear disk space.

Mac OS X Server ships with a number of command-line tools that you can use to monitor disk space on your computer:

- `df`. This command tells you how much space is used and how much is available on every mounted volume.

For example, the following command lists local volumes and displays disk usage:

```
df -Hl
Filesystem      Size  Used Avail Capacity  Mounted on
/dev/disk0s9    40G   38G   2.1G    95%      /
```

In this example, the hard disk is almost full with only 2.1 GB left. This tells you that you should act immediately to free space on your hard disk before it fills up and causes problems for your users.

- `du`. This command tells you how much space is used by specific folders or files. For example, the following command tells you how much space is used by each user's home folder:

```
sudo du -sh /Users/*
3.2M   /Users/Shared
9.3M   /Users/omar
8.8M   /Users/jay
1.6M   /Users/lili
...
```

Knowing who's using most of the space on the hard disk lets you contact users and have them delete unused files.

Note: With Workgroup Manager, you can set disk quotas for users and generate disk usage reports. For more information, see *User Management*.

- `diskspacemonitor`. This command lets you automate the process of monitoring disk space usage. When the amount of free disk space drops below the level you specify, `diskspacemonitor` executes shell scripts that send you a notification. This command defines two action levels:
 - Alert—Sends you a warning message when disk space usage reaches 75%.
 - Recover—Archives rarely used files and deletes unneeded files when disk space usage reaches 85%.

For more information about these commands, see the corresponding man page or *Command-Line Administration*.

Network Monitoring Tools

Degradation in network performance or other network problems can adversely affect the availability of your services. The following network monitoring tools can alert you to problems early, so you can take corrective action to avoid or minimize down time.

- To monitor network activity, use the `tcpdump` utility in Mac OS X Server. This utility prints the headers of incoming and outgoing packets on a network interface that match specified parameters.

Using `tcpdump` to monitor network traffic is especially useful when trying to detect denial of service attacks. For example, the following command monitors incoming traffic on port 80 on your computer:

```
sudo tcpdump -i en0 dst port 80
```

If you detect an unusual number of requests coming from the same source, use Firewall service to block traffic from that source.

For more information about `tcpdump`, see the corresponding man page or *Command-Line Administration*.

- Consider using Ruby, Perl, shell scripts, or AppleScripts to automate the monitoring process. For example, using `tcpdump` to monitor traffic can be time consuming, so automation is necessary.
- Consider using Ethereal, an X11 open source packet sniffing tool that you can run in the X11 environment on Mac OS X Server. Unlike `tcpdump`, this tool has a graphical user interface and a set of powerful network analysis tools.
For more information about Ethereal, see www.ethereal.com/.
- You can use other third-party tools that automatically analyze network traffic and alert you to problems.

Notification in Server Admin

Server Admin has an easy to use notification system that can keep you informed of your server's hard disk or software status. Server Admin will send an email to any address (local or not) when:

- There is less than a certain percentage of free space left on any system hard disk.
- There are Software Update packages available from Apple.

To use the email functionality, the server will start the SMTP (outgoing mail) process on the server. Make sure the firewall allows SMTP traffic from the server.

To set a notification:

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the Notifications tab.
- 3 Click the Add (+) button below the "Addresses to notify" field and add an address.
- 4 Repeat as needed, then click Save.

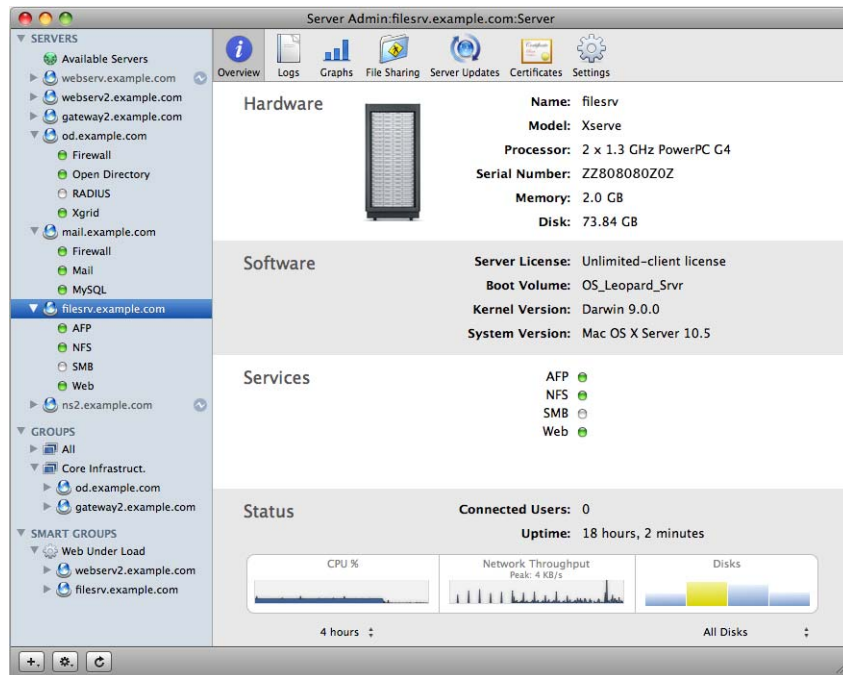
Monitoring Server Status Overviews Using Server Admin

Server Admin has several ways to see a status overview, from detailed information for a single server to a simplified overview for many servers at once.

To see a status overview for one server:

- Select a server in the Server list.

The following shows a sample Overview pane for a single server.

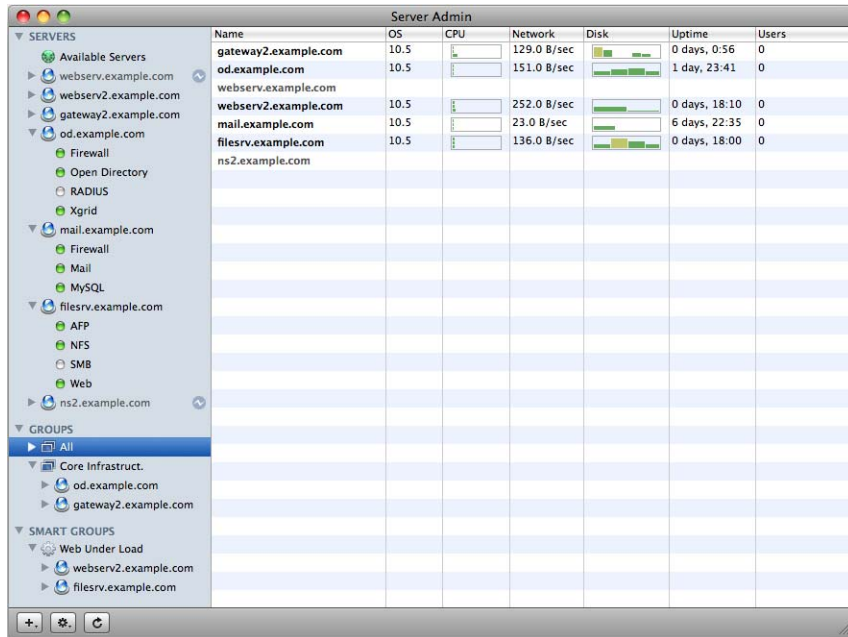


This overview shows basic hardware, operating system versions, active services, and graphs of CPU history, network throughput history, and disk space.

To see status overview of many servers at once:

- Select a server group, smartgroup, All Servers group, or Available Servers group.

The following shows a sample Overview pane for a group of servers.



Name	OS	CPU	Network	Disk	Uptime	Users
gateway2.example.com	10.5		129.0 B/sec		0 days, 0:56	0
od.example.com	10.5		151.0 B/sec		1 day, 23:41	0
websevr2.example.com	10.5		252.0 B/sec		0 days, 18:10	0
mail.example.com	10.5		23.0 B/sec		6 days, 22:35	0
filesrv.example.com	10.5		136.0 B/sec		0 days, 18:00	0
ns2.example.com						

This overview shows the:

- Hostname
- OS version
- Current CPU usage graph (a mouseover reveals more specific numbers)
- Current network throughput
- Disk space used (a mouseover reveals more specific numbers)
- Uptime
- Number of connected file services users

You can sort the list by column.

Simple Network Management Protocol (SNMP)

Simple Network Management Protocol (SNMP) is a common protocol for monitoring the status of network equipment (for example, routers and smart switches), computers, and other networkable devices like Uninterruptable Power Supplies. Mac OS X Server uses Net-SNMP to implement SNMP v1, SNMP v2c, and SNMP v3 using both IPv4 and IPv6.

SNMPv2 is the default access protocol and the default read-only community string is “public.”

Enabling SNMP reporting

SNMP access isn't enabled by default on Mac OS X Server. To use SNMP tools to poll your Mac OS X Server for data you must configure and then enable the service.

To enable SNMP

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the General tab.
- 3 Select Network Management Server (SNMP).
- 4 Click Save.

When SNMP is active, anyone with a route to the SNMP host can collect SNMP data from it.

- 5 Configure the basic SNMP parameters from the command-line.

The SNMP process will not start unless `/etc/snmpd.conf` has been configured for the current site. To configure, see "Configuring snmpd" on page 180.

Note: The default configuration of `snmpd` uses privileged port 161. For this reason and others, it must be executed by root or using `setuid`. You should only use `setuid` as root if you understand the ramifications. If you do not, seek assistance or additional information. Flags available for `snmpd` will change the uid and gid of the process after it starts. For more information, see the `snmpd` man page.

Configuring snmpd

The configuration (`.conf`) file for `snmpd` is typically at `/etc/snmpd.conf`. If you have an environment variable `SNMPCONF`, `snmpd` will read any files named `snmpd.conf` and `snmpd.local.conf` in these directories. The `snmpd` process can be started with a `-c` flag to indicate other conf files. For more information about which conf files can be used, see the `snmpd` man page.

Configuration files can be created and installed more elegantly using the included script `/usr/bin/snmpconf`. As root, use this script with the `-i` flag to install the file at `/usr/share/snmp/`. Otherwise the default location for the file to be written is the user's home folder (`~/`). Only root has write permission for `/usr/share/snmp/`.

Because `snmpd` reads its configuration files at startup, changes to configuration files require that the process be stopped and restarted. You can stop `snmpd` with ProcessViewer or at the command-line (`kill -HUP <pid>`).

To enable and configure SNMP:

- Use the `/usr/bin/snmpconf` command, which takes you through a basic text-based setup assistant for configuring the community name and saves the info in the configuration file.

The `snmp` config file is located in `/usr/share/snmp/snmpd.conf`.

SNMP Configuration Example

Step 1: Customize data

- 1 To customize the data provided by `snmpd`, add an `snmpd.conf` file using `/usr/bin/snmpconf` as root or using `sudo`, by executing this command:

```
/usr/bin/snmpconf -i
```

If there are existing configuration files, you can reading them into the assistant and incorporate their contents with the output of the assistant.

- 2 Choose to read in the file by indicating the file at `/etc/snmp/snmpd.conf`.

You will then see a series of text menus.

- 3 Make these choices in this order:

a Select File: 1 (`snmpd.conf`)

b Select section: 5 (System Information Setup)

c Select section: 1 (The [typically physical] location of the system.)

d The location of the system: type text string here — such as “`server_room`”

e Select section: f (finish)

f Select section: f (finish)

g Select File: q (quit)

You have created an `snmpd.conf` file with a creation date of today.

Verify its creation by entering `ls -l /usr/share/snmpd.conf`.

Step 2: Restart `snmpd` to take changes

- 1 Open Server Admin.
- 2 Select a server, click the Settings button in the toolbar, and then click the General tab.
- 3 Deselect Network Management Server (SNMP).
- 4 Click Save.

You can also do this via the command-line by killing and restarting the `snmpd` process as root:

```
/usr/sbin/snmpd
```

Step 3: Collect SNMP information from the host

- To get the SNMP-available information you just added, execute this command from a host that has SNMP tools installed:

```
/usr/bin/snmpget -c public <hostname> system.sysLocation.0
```

Replace “<hostname>” with the actual name of the target host.

You should see location you provided. In this example, you would see:

```
SNMPv2_MIB::system.sysLocation.0 = STRING:\“server_room\”
```

The other options in the menu you were working in are:

```
/usr/bin/snmpget -c public <hostname> system.sysContact.0  
/usr/bin/snmpget -c public <hostname> system.sysServices.0
```

The final .0 indicates you are looking for the index object. The word public is the name of the snmp community that you did not alter.

If you need information about either of these or if you need explanations of snmp syntax, tutorials are available at net-snmp.sourceforge.net.

Tools to Use with SNMP

Other than snmpget, there are other snmp based tools installed, and third-party suites (both free and commercial) are available with varying complexity and reporting.

Additional Information

Additional information about SNMP can be had from the following sources.

Man pages

Entering `man -k snmp` in the Terminal will provide a list of the known man pages.

Web sites

The Net SNMP-Project:

- www.net-snmp.org
- net-snmp.sourceforge.net

Books

Essential SNMP by Douglas Mauro, Kevin Schmidt

Publisher: O'Reilly (Second Edition Sept 2005)

ISBN: 0-596-00840-6, 460 pages

Notification and Event Monitoring Daemons

To monitor and log system events, the operating system runs several daemons that intercept application messages and log them or act on them.

There are two main notification daemons: syslogd and emond.

- **syslogd:** The syslogd daemon is a standard UNIX method of monitoring systems. It logs messages in accordance with the settings found in `/etc/syslog.conf`. You can examine the output files specified in that configuration by using a file printing or editing utility because they are plain text files. Administrators can edit these settings to fine-tune what is being monitored.

Many administrators will tail or scrape the log file, meaning they will have scripts parse the log files and perform some action if a designated bit of information is present in the log. These home-grown notifications vary in quality and usefulness and are tailored to the script-writer's specific needs.

The syslogd daemon can be configured to send and receive log file information to or from a remote server (by editing the `/System/Library/LaunchDaemons/com.apple.syslogd.plist`). This is not recommended because syslogd does not use secure means to send log messages across the net.

- **emond:** The daemon emond is the event monitoring system for Mac OS X Server v10.5. It is a unified process that handles events passed from other processes, acts on the events as designated in defined rule set, and then notifies the administrator. Currently, emond is the engine used for Server Admin's email notification system. It is not used for Server Monitor's notifications.

The high-level service receives events from the registered client, analyzes whether the event requires handling based on rules provided by the service at the time it registered and, if handling is required, the action related to that event is performed. To accomplish this the daemon emond has three main parts: the rules engine, the events it can respond to, and the actions it can take.

The emond rules engine works in the following manner. It:

- Reads the config info from `/etc/emond.d/emond.conf`.
- Reads in the rules from plist files in the `/etc/emond.d/rules/` directory.
- Processes the startup event.
- Accepts events until terminated.
- Processes the rules associated with the event, triggering as needed.
- Performs actions specified by the rules that were triggered.
- Runs as the least privileged possible (nobody).

WARNING: The file formats and settings in `emond.conf` and rules plists are not documented for customer use. Tampering could result in an unusable notification system and is unsupported.

Logging

Mac OS X Server maintains standard UNIX log files and Apple-specific process logs. Logs for the OS can be found in:

- /var/log
- /Library/Logs
- ~/Library/Logs

Each process is responsible for its own logs, the log level, and verbosity. Each process or application can write its own log file or use a system standard log, like syslog. You can use the Console application (in /Applications/Utilities) to read these and other plain-text log files regardless of location.

Most services in Mac OS X Server have a logging pane in Server Admin. You can use these panes to set logging levels and view the logs for any particular service.

Syslog

The system log, syslog, is a consolidated catch-all location for process log messages. syslog has several levels of available log detail. If low detail logging is selected, detailed messages are not saved, but high detail logging results in large and possibly unhelpfully large log files.

The level of logging you use for syslog can be tuned by process and should be appropriate to the level necessary for successful notification and debugging.

Syslog log levels (in ascending order from least to most detail)

Level name	Level indicator in syslog.conf	Amount of detail
None	.none	None
Emergency	.emerg	Least
Alert	.alert	
Error	.err	
Warning	.warn	
Notice	.notice	
Info	.info	
Debug	.debug	Most

Syslog Configuration File

The configuration file can be found at /etc/syslog.conf. Each line has the following format:

```
<facility>.<loglevel> <path to logfile>
```


Facility is the process name writing to the log, and the path is the standard POSIX path to the log file. Asterisks (*) can be used as wildcards. For example, the setting for the kernel is:

```
kern.* /var/log/system.log
```

This shows that all messages to the log of all levels from the kernel are to be written in the file `/var/log/system.log`.

Likewise, the following setting is an example of all emergency messages from all processes being sent to a custom emergencies log file:

```
*.emerg /var/log/emergencies.log
```

Directory Service Debug Logging

If you are using Open Directory and you want debugging information from Directory Services processes, you must use a different logging method than `systemlog`. You must enable debug logging on the process manually. When enabled, this debug logging writes messages to the log file at:

```
/Library/Logs/DirectoryService/DirectoryService.debug.log
```

The following commands must be performed with superuser permissions (`sudo` or `root`):

To manually turn on/off debug logging for Directory Services:

```
killall -USR1 DirectoryService
```

To start debugging at startup:

```
touch /Library/Preferences/DirectoryService/.DSLogAPIAtStart
```

Note: The debug log is not self-documented and is not intended for normal logging. It is very verbose and very opaque. It shows API calls, plugin queries, and responses.

Open Directory Logging

The configuration file can be found at `/etc/openldap` and the logs are found in `/var/log/slapd.log`. Each directory transaction generates a separate transaction log in the OpenLDAP database. The database and transaction logs can be found at `/var/db/openldap/openldap-data`.

The `slapd` process, which governs Open Directory usage, has an additional parameter for extra logging. The following command enables the additional logging:

```
slapconfig -enablenesslapdlog
```

To run slapd in debugging mode:

- 1 Stop and remove slapd from launchd's watch list:

```
launchctl unload /System/Library/LaunchDaemons/org.openldap.plist
```

- 2 Restart slapd in debug mode:

```
sudo /usr/libexec/slapd -d 99
```

AFP Logging

The server side of Apple File Service Protocol (AFP) keeps track of access and errors, but it does not have much debugging information. However, you can add client-side logging to AFP clients to help monitor and troubleshoot AFP connections.

To enable client-side logging:

Perform all these actions on the AFP client computer.

- 1 Set the client debug level (levels 0-8):

```
defaults write com.apple.AppleShareClientCore -dict-add afp_debug_level 4
```

- 2 Set the client log message recipient (in this case, syslog):

```
defaults write com.apple.AppleShareClientCore -dict-add afp_debug_syslog 1
```

- 3 Enable syslog to catch the debugging messages from the client:

You do this by adding *.debug /var/log/debug.log to the syslog.conf file.

- 4 Restart the syslog process.

Additional Monitoring Aids

You can use additional aids for monitoring Mac OS X Server. There are a number of third-party server monitoring packages, as well as an additional Apple monitoring tool.

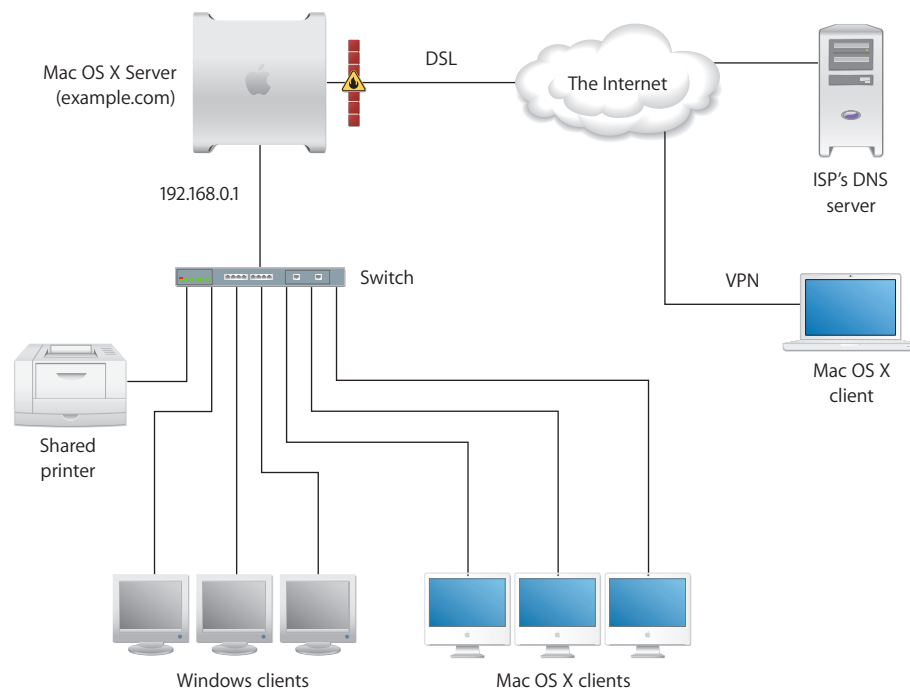
The inclusion of third-party tools in the following list does not constitute an endorsement of or support for these products. They are listed for informational purposes only.

- **Apple Remote Desktop:** This software package contains many features that allow you to interact with, get reports on, and track computers running Mac OS X and Mac OS X Server. It has several powerful administration features and excellent reporting capabilities.
- **Nagios (third-party):** This tool is an open source computer system and network monitoring application.
- **Growl (third-party):** This tool is a centralized, extensible notification service that supports local and remote notification.

The setup example in this chapter illustrates one way to set up the directory and network infrastructure of Mac OS X Server in a small business scenario.

A Single Mac OS X Server in a Small Business

In this example, Mac OS X Server provides directory, network, and productivity services to employees in a small business:



The small business has been using an office LAN to share files and a printer. Acquiring Mac OS X Server made it possible to implement an intranet that uses an ISP's DNS and digital subscriber line (DSL) services.

Here's a summary of the scenario's characteristics:

- An Open Directory master LDAP directory on the server centralizes user management, including authentication of Mac OS X and Windows users.
- The ISP's DNS service provides a DNS domain name for the company (example.com).
- A DNS server running on Mac OS X Server provides name services for the server, the printer, and any other intranet device that has a static IP address.
- A firewall between the server and the Internet protects the intranet from unauthorized access.
- NAT service lets intranet users share the ISP's IP address for Internet access, while VPN lets employees access the intranet securely over the Internet when employees work away from the office.
- DHCP service on Mac OS X Server provides dynamic IP addresses to intranet client computers. The server and printer have static addresses, but client computers have dynamic addresses.

How to Set Up the Server

The following steps summarize how to set up Mac OS X Server in this hypothetical small business. For complete information about setting up directory services, see *Open Directory Administration*. For details about network service setup (IP firewall, DHCP, and so forth), see *Network Services Administration*.

Step 1: Set up the network

- 1 Make sure the server has two Ethernet interfaces (ports): one for the intranet (LAN) connection and one for the DSL modem connection.

Use the faster interface for the server connection. A 10-Mbit connection is more than sufficient for the DSL connection.

- 2 Connect the server to the LAN using the faster interface.

In this example, the server is plugged in to a switch used to connect client computers and shared printer. We'll refer to this interface as the internal interface.

Intranet devices should be connected to a hub or switch using good-quality CAT-5 Ethernet cables. A high-speed 10/100/1000 megabit switch can support advanced server features such as NetBoot that work best over a fast connection.

- 3 Connect the server to the DSL modem using the other Ethernet interface.

We'll refer to this interface as the external interface.

Step 2: Contact the ISP to set up external DNS

The ISP's Name Servers should be serving the company zone example.com containing all public IPs of all servers and services available to the Internet (for example, the company web server and the VPN gateway).

This means that the zone handled by the ISP contains only the public IP addresses and the ISP's name server provides the necessary redundancy. The ISP should also provide Forward and Reverse DNS lookup for the zone's domain for any external IP Address being used.

WARNING: This example assumes that the ISP is providing Forward and Reverse DNS resolution for the public IP address and machine name of the server. If this is not the case (for example, if your ISP's setup is not done yet or you plan to run your own name server on the server itself), choose Standalone Server in Step 4 and promote it to an Open Directory Master or Replica only after there is a working DNS setup.

Step 3: Set up an administration computer

- 1 Install the server administration tools from the Server Tools DVD.

Choose a computer running Mac OS X Leopard to install the tools on. Make sure the network communication between the administrator computer and the target server is functioning. For more instructions, see "Preparing an Administrator Computer" on page 82.

- 2 Fill out the "Mac OS X Server Advanced Worksheet" in the appendix on page 197. You'll need the information as you move through the Assistant's panes.

Step 4: Set up the server and the master directory

- 1 Start the server from the Install DVD.

The procedure you use depends on the server hardware.

In this example, assume the computer has a keyboard and a DVD drive. Turn on the computer, insert the Install DVD into the optical drive, and restart the computer while holding down the C key on the keyboard.

Chapter 5, "Installation and Deployment," on page 79 has instructions for other installation methods, such as installing on a server without an optical drive and installing from a NetInstall environment.

- 2 Start up Setup Assistant on the administrator computer.
- 3 When the Setup Assistant opens, choose "Install Mac OS X Server on a remote computer."

- 4 Proceed by following the onscreen instructions.
If you need to format the target disk, see “Preparing Disks for Installing Mac OS X Server” on page 91 for instructions on preparing disks for installing Mac OS X Server.
When installation is complete, the server restarts.
- 5 After restarting, use Server Assistant again and choose “Set up a remote computer.”
- 6 Use the Language and Keyboard panes to reflect the server’s administration language.
- 7 In the Administrator Account pane, enter the server administrator’s names and password, and then click Continue.
- 8 In the Network Names pane, if you don’t see the newly installed server, click the Add (+) button, enter the IP address, and enter the default administrator name and password, and click Continue.
For more information, see “Connecting to the Network During Initial Server Setup” on page 108.
- 9 Proceed by following the onscreen instructions.
- 10 Make sure the Network Interfaces pane lists external and internal Ethernet interfaces.
- 11 Make sure the external interface is the first one listed in the Network Interfaces pane.
The first interface listed is the primary, or default, interface. Network traffic initiated by the server is routed through the primary interface. VPN uses it as the Public network, treating all others listed as Private.
- 12 Click Continue.
The TCP/IP Connection pane appears for each Ethernet interface.
- 13 For the external interface, choose Manually from the Configure IPv4 pop-up list, then enter the IP address, subnet mask, and DNS server IP address or addresses provided to you by the ISP.
With a dual interface setup like the one in this example, all DNS requests are routed to the primary interface. So when running DNS on your server, enter the gateway’s public IP in the Name Servers field as well. In a manual configuration, make it appear first in the list so it is consulted before your ISP’s servers, then click Continue.
- 14 If you’ll be using Gateway Setup Assistant (from the NAT service section of Server Admin) to configure network settings, you don’t need to set up an internal interface. Otherwise, enter these values for the internal interface then click Continue:
 - Configure IPv4: Manually
 - IP Address: 192.168.0.1 (192.168 values are reserved for internal LANs)
 - Subnet Mask: 255.255.0.0
 - Router: 192.168.0.1
 - DNS servers: 192.168.0.1

- 15 In the Directory Usage Pane, choose Open Directory Master to set up a shared LDAP directory on the server; then Select Enable Windows Primary Domain Controller and enter a Domain/Workgroup name.

These settings will set up a Windows PDC so that employees who use Windows NT, Windows 2000, and Windows XP workstations can log in to the PDC, change passwords during login, and have roaming user profiles and network home folders on the server.

With one user account, a user can log in from a Windows workstation or a Mac OS X computer and access the same network home folder.

- 16 Click Continue.
- 17 Proceed through the remaining Assistant panes, then click Apply to initiate server setup.

When setup is complete, the server restarts.

- 18 Log in to the server as the administrator you defined when using Server Assistant.
- 19 Configure the server's network settings.

The simplest way to do this is to use the Gateway Setup Assistant, as Step 4 describes. Alternatively, you can individually configure each network service using Server Admin, as Steps 5 through 8 describe.

Step 5: Use Gateway Setup Assistant to automate the server's network configuration

- 1 Open Server Admin on the administrator computer.
- 2 If you have not already done so, connect and authenticate to the server as the administrator you defined when using Server Assistant.
- 3 Select the server and add the services you are going to use.
For this step, select NAT service and Firewall service.
- 4 In the Overview pane of the server you're setting up, click on the NAT service.
- 5 Open Gateway Setup Assistant by clicking the button on the NAT overview pane.
- 6 Proceed through the panes, specifying information when prompted.
On the WAN Port pane, select the port you configured during initial setup as the external interface.
On the VPN settings pane, enable VPN and specify a shared secret for client connections to use.
On the LAN Ports pane, select the port you want to use as the internal interface.
- 7 When Gateway Setup Assistant has completed network setup and you've quit the application, go to Step 9.

Step 6: Set up the firewall

- 1 Open Server Admin on the administrator computer.
- 2 If you have not already done so, connect and authenticate to the server as the administrator you defined when using Server Assistant.
- 3 In the service list, click Firewall.
- 4 Click Start Firewall in the bottom action bar.
- 5 Click Settings and select Services.
- 6 Choose Edit Services for the address group named "192.168-net."
- 7 Select "Allow" for services you want employees working at the office to be able to access.
At a minimum, select Domain Name Service, DHCP, and NetBoot.
- 8 Choose to Edit Services for the address group named "any."
- 9 Click Services and select Allow for services you want external clients to be able to access behind the firewall. At a minimum, select L2TP VPN, IKE, and DHCP.
- 10 Click Save.

Step 7: Set up DNS service

The DNS of Leopard Server handles zone information (for example, all fully qualified host names for the local site like "site1.example.com"), mapping this private zone to private, local IPs. This avoids the need to add public servers to the local DNS.

Additionally, a DNS forwarder zone is set up to query the ISP's DNS records for anything not found in the local DNS zone (for example, the IP addresses of other organization's web servers like www.apple.com).

Note: As noted in Step 2 this example assumes that your ISP is providing Forward and Reverse DNS for your company's zone <example.com>, including resolution of the server's public IP.

As a result, the inhouse name server uses an internal zone like <site1.example.com>, which holds the private IP addresses of the server and all other devices on the LAN.

- 1 In Server Admin, select DNS in the service list.
- 2 Click Zones, click the Add button (+) under the Zones list, and select Add Primary Zone.
- 3 Select the default zone, and customize it to fit your organization.

In this case, settings are:

- Primary Zone Name: example.com
- Nameservers Address: 192.168.0.1
- Administrator email: admin@example.com

- 4 Add a machine record to the zone, by selecting the zone, clicking “Add Record,” and selecting “Add Machine (A)” from the pop-up button.
- 5 Using the following settings, select the machine record which is under the zone name to edit the record, and clicking Save when finished.
 - Machine name: myserver
 - IP Address: 192.168.0.1
- 6 Using the following settings, continue to add machines to the zone.

For example, to add a printer, click the Add button, specify values for the printer, then click OK:

 - IP address: 192.168.100.2
 - Name: laserprinter_2000
- 7 Set the server to look outside the server for any domain name it doesn’t control, by clicking Settings.
- 8 In the Forwarder IP Addresses list, click the Add (+) button to add the DNS addresses provided by the ISP.
- 9 Click Save, then click Start DNS.

Step 8: Set up DHCP service

This step sets up a DHCP server that provides employee computers with dynamic IP addresses as well as the identity of the DNS, LDAP, and WINS servers they should use.

When a client computer’s search policy is set to Automatic (using the Directory Utility application on the client computer), the identity of the DNS, LDAP, and WINS servers is supplied when an IP address is supplied.

- 1 In Server Admin, make sure DNS is running.
- 2 Select DHCP in the service list.
- 3 Click Subnets.
- 4 Click the Add (+) button to define the range of addresses to dynamically assign.

The range should be large enough to accommodate current and future client computers. Make sure you exclude some addresses (at the start or end of the range) so they’re reserved for devices that need static IP addresses or for VPN users.

Here are some sample values:

 - Subnet Mask: 255.255.0.0
 - Starting IP Address: 192.168.0.2
 - Ending IP Address: 192.168.0.102
 - Network Interface: en1
 - Router: 192.168.0.1

- 5 Make sure the DNS pane contains the following values:
 - Default Domain: example.com
 - Name Servers: 192.168.0.1
- 6 Click LDAP to configure DHCP to identify the server you're configuring as the source of directory information for clients who are served dynamic IP addresses.

The server you're setting up should be identified in the Server Name field because you set up the server as an Open Directory master when you used Server Assistant. Other settings are optional for this example.
- 7 Click WINS to configure DHCP to serve Windows-specific settings to clients who are served dynamic IP addresses; then supply these values:
 - WINS/NBNS Primary Server: 192.168.0.1
 - NBT Node Type: Broadcast (b-node)
- 8 Click Save, enable the internal Ethernet interface, then click Start DHCP.

Step 9: Set up NAT service

- 1 In Server Admin, select NAT in the service list.
- 2 Click Settings.
- 3 Select the external interface from the "External network interface" pop-up menu.
- 4 Click Save, then click Start NAT.

Step 10: Set up VPN service

- 1 In Server Admin, select VPN in the service list.
- 2 Click Settings.
- 3 Enable L2TP over IPSec (Layer Two Tunneling Protocol, Secure Internet Protocol) for Mac OS X v10.5 computer users, Linux or UNIX workstation users, and Windows XP users.

Although PPTP can also be used, L2TP provides the greatest security because it runs over IPSec.
- 4 Enter a starting and ending IP address to indicate the addresses the VPN server can assign to clients.

Avoid addresses the DHCP server is set up to serve. Also avoid addresses you specify if you enable PPTP.
- 5 Specify the shared secret by entering a string in "Shared secret" that isn't intuitive.

For example, specify digits, symbols, and uppercase and lowercase characters in unusual combinations. The recommended length is 8 to 12 characters.

- 6 Enable Point to Point Tunneling Protocol (PPTP) if employees will need to access the intranet from Windows workstations other than Windows XP computers or from Mac OS X v10.2 computers when they're away from the office.

If you need to support older Windows clients that don't have 128-bit PPTP support, select "Allow 40-bit encryption keys in addition to 128-bit."

- 7 Enter a starting and ending IP address to indicate the addresses the VPN server can assign to clients.

Avoid addresses the DHCP server is set up to serve. Also avoid addresses you specified when you enabled L2TP over IPSec.

- 8 Click Save, then click Start VPN.

Step 11: Set up productivity services

The infrastructure you need to set up file, print, and other productivity services is now available. Follow the instructions in the relevant administration guides, listed on page 13, to configure the services of interest.

Many services, such as Apple File service, require minimal setup. Simply start them using Server Admin.

Step 12: Create user accounts and home folders

- 1 Open Workgroup Manager.
- 2 If you have not already done so, connect and authenticate to the server as the administrator you defined when using Server Assistant.

The Open Directory master LDAP directory is available for editing. You'll add an account for each employee to this master directory.

- 3 Click the New User button.
- 4 Specify user settings in the panes that appear.

User Management tells you how to set up all user account attributes, including home folders. It also describes how to manage users by setting up group accounts and computer lists and how to set up preference settings that customize the work environments of Macintosh clients.

User Management and *Open Directory Administration* show how to implement support specifically for Windows workstation users.

Step 13: Configure client computers

The information that follows applies to Mac OS X v10.5 computers.

- 1 If necessary, configure Mac OS X clients to retrieve information from the DHCP server.

Mac OS X v10.5 computers are configured to use DHCP to obtain IP addresses and retrieve information about an LDAP directory from the DHCP server. After you configure DHCP service with information about an LDAP directory, that information is delivered to Mac OS X clients when they receive IP addresses from the DHCP server.

These settings are preconfigured:

- Network preferences are set to use DHCP. To access the setting, select System Preferences, open Network preferences, select the internal Ethernet interface, and select “Using DHCP with manual address” or “Using DHCP” from the Configure IPv4 pop-up menu.
- The computer’s search policy is set to be defined automatically. To access this setting, open Directory Utility (in /Applications/Utilities/) and click Authentication. If the lock icon is locked, click it and authenticate as an administrator. Choose Automatic from the Search pop-up menu, then click Apply.
- The use of DHCP-supplied LDAP information is enabled. To access this setting, open Directory Utility and click Services. If the lock icon is locked, click it and authenticate as an administrator. Select LDAPv3 in the list of services, then click Configure. Click “Use DHCP-supplied LDAP Server,” then click OK.

- 2 Configure Mac OS X clients so they can use the VPN server.
- 3 Open the Internet Connect application (in /Applications/) and click VPN in the toolbar.
- 4 Select L2TP over IPSec or PPP and click Continue.
- 5 From the Configurations pop-up menu., choose Edit Configurations
- 6 Enter the external IP address from the ISP, the user name and password for the computer user and, for L2TP over IPSec, the shared secret.
- 7 Click OK.

Mac OS X Server Advanced Worksheet

Enter settings for the server in the tables below:

Server name:

Item	Description	Your information
Identity of remote server for installation and setup	<p>For interactive installation and setup of a remote server on the local subnet, one of these values for the server:</p> <ul style="list-style-type: none"> - IP address in IPv4 format (000.000.000.000) - host name (someserver.example.com) - MAC address (00:03:93:71:26:52). <p>For command-line or remote-subnet installations and setups, the target server's IP address, in IPv4 format.</p>	
Preset password (for remote installation and setup)	<p>The first eight digits of the target server's built-in hardware serial number, printed on a label on the computer.</p> <p>For older computers with no such number, use 12345678 for the password.</p>	
Type of installation	<p>Upgrade from the latest 10.4 version or from v10.3.9, complete installation without disk formatting, or clean installation.</p> <p>The target volume (partition) is erased when you do a clean installation.</p>	
Target disk or partition	Name of the target disk or partition (volume).	
Disk format (when erasing the disk is OK)	<p>A format for the target disk.</p> <p>In most cases, use Mac OS Extended (Journaled). You can also use Mac OS Extended or case-sensitive HFS+.</p>	
Disk partitioning (when erasing the disk is OK)	<p>Indicate whether you want to partition the target disk.</p> <p>The minimum recommended size of a target disk partition is 4 GB.</p>	

Item	Description	Your information
RAID mirroring (when erasing the disk is OK and you have a second physical drive on the target server)	<p>Indicate whether you want to set up RAID mirroring. The second disk is used automatically if the primary disk isn't available.</p> <p>If the target disk has a single partition and the second physical drive has a single partition and no data, you can set up RAID mirroring after installation. However, to prevent data loss, set up RAID mirroring as soon as possible.</p>	
Using saved setup data	<p>If you want to use saved setup data to set up this server, identify the file or directory storing the data you want to use. If the data is encrypted, also identify the passphrase.</p> <p>If you want to save settings in a file or directory, use one of the next two rows.</p>	
Saving setup data in a file	<p>Name the file using one of these options:</p> <ul style="list-style-type: none"> • <MAC-address-of-server>.plist (include leading zeros but omit colons, for example, 0030654dbcef.plist). • <IP-address-of-server>.plist (for example, 10.0.0.4.plist). • <partial-DNS-name-of-server>.plist (for example, myserver.plist). • <built-in-hardware-serial-number-of-server>.plist (first eight characters, for example, ABCD1234.plist). • <fully-qualified-DNS-name-of-server>.plist (for example, myserver.example.com.plist). • <partial-IP-address-of-server>.plist (for example, 10.0.plist matches 10.0.0.4 and 10.0.1.2). • generic.plist (a file that any server will recognize, used to set up servers that need the same setup values) <p>If you encrypt the file, you can save the passphrase in a file named using the above conventions, except use the extension .pass, not .plist.</p> <p>Place the files in a location where the target server or servers can detect it. A server can detect files that reside on a volume mounted locally in /Volumes/*/Auto Server Setup/, where * is any device mounted under /Volumes.</p>	

Item	Description	Your information
Saving setup data in a directory	<p>Navigate to the directory where you want to save the setup, and name the setup record using one of these options:</p> <ul style="list-style-type: none"> • <MAC-address-of-server> (include leading zeros but omit colons, for example, 0030654dbcef). • <IP-address-of-server> (for example, 10.0.0.4). • <partial-DNS-name-of-server> (for example, myserver). • <built-in-hardware-serial-number-of-server> (first eight characters, for example, ABCD1234). • <fully-qualified-DNS-name-of-server> (for example, myserver.example.com). • <partial-IP-address-of-server> (for example, 10.0 matches 10.0.0.4 and 10.0.1.2). • generic (a record that any server will recognize, used to set up servers that need the same setup values) <p>If you encrypt the file, you can save the passphrase in a file named using the above conventions, except add the extension .pass. Place the passphrase file in a location where the target server or servers can detect it. A server can detect the file if it resides on a volume mounted locally in /Volumes/*/Auto Server Setup/, where * is any device mounted under /Volumes.</p>	
Language	<p>The language to use for server administration (English, Japanese, French, or German). The language affects the server's time and date formats, displayed text, and the default encoding used by the AFP server.</p>	
Keyboard layout	<p>The keyboard for server administration.</p>	

Item	Description	Your information
Serial number	<p>The serial number for your copy of Mac OS X Server. You need a new serial number for Mac OS X Server v10.5.</p> <p>The format is xsvr-999-999-x-xxx-xxx-xxx-xxx-xxx-x, where x is a letter and 9 is a digit. The first element (xsvr) and the fourth one (x) must be lower case.</p> <p>Unless you have a site license, you need a unique serial number for each server. You'll find the server software serial number printed on the materials provided with the server software package.</p> <p>If you have a site license, you must enter the registered owner name and organization as specified by your Apple representative.</p> <p>If you set up a server using a generic setup file or directory record and the serial number isn't site-licensed, you must enter the server's serial number using Server Admin.</p>	
Administrator's long name (sometimes called full name or real name)	<p>A long name can contain no more than 255 bytes. The number of characters ranges from 255 Roman characters to as few as 85 3-byte characters.</p> <p>It can include spaces.</p> <p>It can't be the same as any predefined user name, such as System Administrator. This name is case sensitive in the login window, but not when accessing file servers.</p>	
Administrator's short name	<p>A short name can contain as many as 255 Roman characters, typically eight or fewer.</p> <p>Use only a through z, A through Z, 0 through 9, _ (underscore), or - (hyphen).</p> <p>Avoid short names that Apple assigns to predefined users, such as "root."</p>	
Administrator's password	<p>This value is case sensitive and must contain at least 4 characters. It is also the password for the root user.</p> <p>If you record this value, be sure to keep this worksheet in a safe place.</p> <p>After setup, use Workgroup Manager to change the password for this account.</p>	

Item	Description	Your information
Host name	<p>You can't specify this name during server setup. Server Assistant sets the host name to AUTOMATIC in /etc/hostconfig.</p> <p>This setting causes the server's host name to be the first name that's true in this list:</p> <ul style="list-style-type: none"> - The name provided by the DHCP or BootP server for the primary IP address - The first name returned by a reverse DNS (address-to-name) query for the primary IP address - The local hostname - The name "localhost" 	
Computer name	<p>The AppleTalk name and the default name used for SLP/DA. Specify a name 63 characters or fewer but avoid using =, :, or @.</p> <p>The Network browser in the Finder uses SMB to find computers that provide Windows file sharing. Spaces are removed from a computer name for use with SMB, and the name can contain no more than 15 characters, no special characters, and no punctuation.</p>	
Local hostname	<p>The name that designates a computer on a local subnet.</p> <p>It can contain lowercase letters, numbers, and/or hyphens (but not at the ends). The name ends with ".local" and must be unique on a local subnet.</p>	
Network interface data	<p>Your server has a built-in Ethernet port and can have an additional Ethernet port built in or added on. Record information for each port you want to activate.</p>	<p>Use the table provided later in this worksheet to record data for each port.</p>
Directory usage	<p>Select one:</p> <ul style="list-style-type: none"> - Standalone Server (use only the local directory). - Connected to a Directory System (get information from another server's shared directory). If you choose this option, use one of the next four rows in this table to indicate how the server will connect with the directory. - Open Directory Master (provide directory information to other computers). If you choose this option, use the row for "Using Open Directory Master." - No change (for upgrades only). 	
Using "As Specified by DHCP Server"	<p>The directory to use is identified by a DHCP server set up to provide the address and search base of an LDAP server (DHCP option 95).</p>	

Item	Description	Your information
Using "Open Directory Server"	The directory to use is an LDAP directory identified by a DHCP server or identified by specifying an IP address or domain name for the LDAP server.	
Using "Other Directory Server"	The directories to use is configured using the Directory Utility application after you finish setting up the server.	
Using "Open Directory Master"	Optionally indicate if you want to enable a Windows Primary Domain Controller on the server. Provide a Windows computer name and domain for the server. The computer name and domain can contain a-z, A-Z, 0-9, -, but no . or space and can't contain only numbers. Finish setting up the directory you want to host by using Server Admin after completing server setup.	
Time zone	Choose the time zone you want the server to use.	
Network time	Optionally indicate a Network Time Server for the server. Apple recommends that you keep your server's clock accurate by synchronizing it with a network time server.	

Configuration settings for the following port appear in the table below:

Port Name: Built-in Ethernet

Item	Description	Your information
Device name	A UNIX name for the port in the format enx, where x starts with 0. For the value of x for the port you're describing, see your hardware manual. The value en0 always designates a built-in Ethernet port.	en0
Ethernet address	The Media Access Control (MAC) address of the port (00:00:00:00:00:00). This value is usually on a sticker on the server hardware, but you can run Apple System Profiler or a command-line tool such as networksetup to discover the value.	
TCP/IP and AppleTalk	Indicate whether you want to enable the port for TCP/IP and/or AppleTalk. You can connect a port to the Internet by enabling TCP/IP and use the same or a different port for AppleTalk. Enable no more than one port for AppleTalk.	
Order of ports	If you enable more than one port, indicate the order in which the ports should be accessed when trying to connect to a network. All nonlocal network traffic uses the first active port.	

Item	Description	Your information
TCP/IP settings	Use one of the next four rows in this table.	
"Manually"	<p>Specify these settings to manually specify TCP/IP settings:</p> <ul style="list-style-type: none"> - IP address (000.000.000.000). A unique static address. - Subnet mask (000.000.000.000). Used to locate the subnet on the local area network where the server resides. This mask is used to derive the network part of the server's address. What remains identifies the server computer on that network. - Router (000.000.000.000) that supports the subnet the server's on. The router is the machine on the local subnet that messages are sent to the target IP address isn't on the local subnet. - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified DNS names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	
"Using DHCP with Manual IP address"	<p>Specify these settings to use a DHCP server to assign a static IP address and optionally other settings for the port.</p> <p>Make sure the DHCP server is set up and DHCP service running when you initiate server setup:</p> <ul style="list-style-type: none"> - IP address (000.000.000.000). A unique static address. - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified DNS names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	

Item	Description	Your information
"Using DHCP"	<p>Specify these settings if you want to use a DHCP server to assign a dynamic IP address and optionally other settings for the port. Make sure the DHCP server is set up and DHCP service running when you initiate server setup:</p> <ul style="list-style-type: none"> - DHCP client ID (optional). A string that's useful for recognizing a port when its IP address changes. Don't specify a DHCP client ID when using Server Assistant to set up the server remotely. Instead, after setup, use the server's Network preferences to define a DHCP client ID. - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified DNS names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	
"Using BootP"	<p>Specify these settings if you want to use a Bootstrap Protocol server to assign an IP address for the identified port.</p> <p>With BootP, the same IP address is always assigned to a particular network interface. It's used primarily for computers that start up from a NetBoot image:</p> <ul style="list-style-type: none"> - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified domain names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	

Item	Description	Your information
IPv6	<p>To configure IPv6 addressing for the port, select Automatically or Manually.</p> <p>Choose Automatically if you want the server to automatically generate an IPv6 address for the port.</p> <p>Choose Manually to specify IPv6 settings:</p> <ul style="list-style-type: none"> - IPv6 address. Generally written in the form 0000:0000:0000:0000:0000:0000:0000. - Router. The IPv6 address of the router on the local subnet. - Prefix length. The number of significant bits in the subnet mask that are used to identify the network. 	
Ethernet settings	<p>To automatically configure Ethernet settings for the port, choose Automatically.</p> <p>Choose Manually (Advanced) to specify settings if you have specific requirements for the network the server's connected to. Incorrect Ethernet settings can affect network performance or render a port unusable:</p> <ul style="list-style-type: none"> - Speed. The maximum Ethernet speed, in number of bits per second, that can be transmitted using the port. Select one of these options: autoselect, 10baseT/UTP, 100baseTX, and 1000baseTX. - Duplex. Determine whether input and output packets are transmitted at the same time (full-duplex) or alternately (half-duplex). - Maximum Packet Transfer Unit Size (MTU). The largest packet the port will send or receive, expressed in bytes. Increasing the packet size improves throughput, but the devices that receive the packet (switches, routers, and so forth) must support the packet size. Select one of these options: Standard (1500), Jumbo (9000), or Custom (enter a value from 72 to 1500). 	

Configuration settings for the following port appear in the table below:

Port Name:

Item	Description	Your information
Device name	A UNIX name for the port in the format enx, where x starts with 0. For the value of x for the port you're describing, see your hardware manual. The value en0 always designates a built-in Ethernet port.	
Ethernet address	The Media Access Control (MAC) address of the port (00:00:00:00:00:00). This value is usually on a sticker on the server hardware, but you can run Apple System Profiler or a command-line tool such as networksetup to discover the value.	

Item	Description	Your information
TCP/IP and AppleTalk	<p>Indicate whether you want to enable the port for TCP/IP and/or AppleTalk.</p> <p>You can connect a port to the Internet by enabling TCP/IP and use the same or a different port for AppleTalk. Enable no more than one port for AppleTalk.</p>	
Order of ports	<p>If you enable more than one port, indicate the order in which the ports should be accessed when trying to connect to a network. All nonlocal network traffic uses the first active port.</p>	
TCP/IP settings	<p>Use one of the next four rows in this table.</p>	
"Manually"	<p>Specify these settings to manually specify TCP/IP settings:</p> <ul style="list-style-type: none"> - IP address (000.000.000.000). A unique static address. - Subnet mask (000.000.000.000). Used to locate the subnet on the local area network where the server resides. This mask is used to derive the network part of the server's address. What remains identifies the server computer on that network. - Router (000.000.000.000) that supports the subnet the server's on. The router is the machine on the local subnet that messages are sent to the target IP address isn't on the local subnet. - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified DNS names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	
"Using DHCP with Manual IP address"	<p>Specify these settings to use a DHCP server to assign a static IP address and optionally other settings for the port.</p> <p>Make sure the DHCP server is set up and DHCP service running when you initiate server setup:</p> <ul style="list-style-type: none"> - IP address (000.000.000.000). A unique static address. - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified DNS names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	

Item	Description	Your information
"Using DHCP"	<p>Specify these settings if you want to use a DHCP server to assign a dynamic IP address and optionally other settings for the port. Make sure the DHCP server is set up and DHCP service running when you initiate server setup:</p> <ul style="list-style-type: none"> - DHCP client ID (optional). A string that's useful for recognizing a port when its IP address changes. Don't specify a DHCP client ID when using Server Assistant to set up the server remotely. Instead, after setup, use the server's Network preferences to define a DHCP client ID. - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified DNS names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	
"Using BootP"	<p>Specify these settings if you want to use a Bootstrap Protocol server to assign an IP address for the identified port.</p> <p>With BootP, the same IP address is always assigned to a particular network interface. It's used primarily for computers that start up from a NetBoot image:</p> <ul style="list-style-type: none"> - DNS servers (000.000.000.000) used to convert IP addresses to fully qualified domain names and vice versa for the port. - Search domains (optional). Names to automatically append to Internet addresses when you don't fully qualify them. For example, if you specify campus.univ.edu as a search domain, you can enter server1 in the Finder's Connect To Server dialog box to connect to server1.campus.univ.edu. 	

Item	Description	Your information
IPv6	<p>To configure IPv6 addressing for the port, select Automatically or Manually.</p> <p>Choose Automatically if you want the server to automatically generate an IPv6 address for the port.</p> <p>Choose Manually to specify IPv6 settings:</p> <ul style="list-style-type: none"> - IPv6 address. Generally written in the form 0000:0000:0000:0000:0000:0000:0000. - Router. The IPv6 address of the router on the local subnet. - Prefix length. The number of significant bits in the subnet mask that are used to identify the network. 	
Ethernet settings	<p>To automatically configure Ethernet settings for the port, choose Automatically.</p> <p>Choose Manually (Advanced) to specify settings if you have specific requirements for the network the server's connected to. Incorrect Ethernet settings can affect network performance or render a port unusable:</p> <ul style="list-style-type: none"> - Speed. The maximum Ethernet speed, in number of bits per second, that can be transmitted using the port. Select one of these options: autoselect, 10baseT/UTP, 100baseTX, and 1000baseTX. - Duplex. Determine whether input and output packets are transmitted at the same time (full-duplex) or alternately (half-duplex). - Maximum Packet Transfer Unit Size (MTU). The largest packet the port will send or receive, expressed in bytes. Increasing the packet size improves throughput, but the devices that receive the packet (switches, routers, and so forth) must support the packet size. Select one of these options: Standard (1500), Jumbo (9000), or Custom (enter a value from 72 to 1500). 	

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 to share files and network services. AFP uses TCP/IP and other protocols to support communication between computers on a network.

alphanumeric Containing characters that include letters, numbers, and punctuation characters (such as _ and ?).

Apache An open source HTTP server integrated into Mac OS X Server. You can find detailed information about Apache at www.apache.org.

application server Software that runs and manages other applications, usually web applications, that are accessed using a web browser. The managed applications reside on the same computer where the application server runs.

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.

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.

back up (verb) The act of creating a backup.

backup (noun) A collection of data that’s stored for the purpose of recovery in case the original copy of data is lost or becomes inaccessible.

bandwidth The capacity of a network connection, measured in bits or bytes per second, for carrying data.

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.

blog A webpage that presents chronologically ordered entries. Often used as an electronic journal or newsletter.

boot ROM Low-level instructions used by a computer in the first stages of starting up.

BSD Berkeley Software Distribution. A version of UNIX on which Mac OS X software is based.

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.

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 of 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 **certificate, public key infrastructure**.

certification authority See Certificate Authority.

cleartext Data that hasn't been encrypted.

client A computer (or a user of the computer) that requests data or services from another computer, or server.

command line The text you type at a shell prompt when using a command-line interface.

command-line interface A way of interacting with the computer (for example, to run programs or modify file system permissions) by entering text commands at a shell prompt. See also **shell**.

computer list A set of computers that all receive the managed preference settings defined for the list, and that are all available to a particular set of users and groups. A computer can be a member of only one computer list. Computer lists are created in Mac OS X Server version 10.4 or earlier.

computer name The default name used for SLP and SMB service registrations. The Network Browser in the Finder uses SLP to find computers advertising Personal File Sharing and Windows File Sharing. It can be set to bridge subnets depending on the network router settings. When you turn on Personal File Sharing, users see the computer name in the Connect to Server dialog in the Finder. Initially it is “<first created user>’s Computer” (for example, “John’s Computer”) but can be changed to anything. The computer name is used for browsing for network file servers, print queues, Bluetooth® discovery, Apple Remote Desktop clients, and any other network resource that identifies computers by computer name rather than network address. The computer name is also the basis for the default local host name.

CUPS Common UNIX Printing System. A cross-platform printing facility based on the Internet Printing Protocol (IPP). The Mac OS X Print Center, its underlying print system, and the Mac OS X Server print service are based on CUPS. For more information, visit www.cups.org.

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**.

default The automatic action performed by a program unless the user chooses otherwise.

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.

DHCP lease time See **lease period**.

digital signature An electronic signature that can be used to verify the identity of the sender of a message.

directory domain A specialized database that stores authoritative information about users and network resources; the information is needed by system software and applications. The database is optimized to handle many requests for information and to find and retrieve information quickly. Also called a directory node or simply a directory.

directory node See **directory domain**.

directory services Services that provide system software and applications with uniform access to directory domains and other sources of information about users and resources.

disc Optical storage media, such as a CD or DVD.

disk A rewritable data storage device. See also **disk drive**.

disk drive A device that contains a disk and reads and writes data to the disk.

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.

DNS domain A unique name of a computer used in the Domain Name System to translate IP addresses and names. Also called a **domain name**.

DNS name A unique name of a computer used in the Domain Name System to translate IP addresses and names. Also called a **domain 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."

domain name See **DNS name**.

Domain Name System See **DNS**.

DSL Digital subscriber line. A broadband data transmission technology that operates over telephone lines.

Dynamic Host Configuration Protocol See **DHCP**.

dynamic IP address An IP address that's assigned for a limited period of time or until the client computer no longer needs it.

EFI Extensible Firmware Interface. Software that runs automatically when an Intel-based Macintosh first starts up. It determines the computer's 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.

Ethernet ID See **MAC address**.

everyone Any user who can log in to a file server: a registered user or guest, an anonymous FTP user, or a website visitor.

export In the Network File System (NFS), a way of sharing a folder with clients on a network.

failover In Xsan, the automatic process by which a standby metadata controller becomes the active metadata controller if the primary controller fails.

Fast Ethernet A group of Ethernet standards in which data is transmitted at 100 megabits per second (Mbit/s).

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.

file system A scheme for storing data on storage devices that allows applications to read and write files without having to deal with lower-level details.

filter A screening method to control access to a server. A filter is made up of an IP address and a subnet mask, and sometimes a port number and access type. The IP address and the subnet mask determine the range of IP addresses that the filter applies to.

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.

FireWire A hardware technology for exchanging data with peripheral devices, defined by IEEE Standard 1394.

format (verb) In general, to prepare a disk for use by a particular file system.

forward zone The DNS zone that holds no records of its own, but forwards DNS queries to another zone.

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.

gateway A network node that interfaces one network to another. Often, it refers to a computer that links a private LAN to a public WAN, with or without Network Address Translation (NAT). A router is a special kind of gateway that links related network segments.

GB Gigabyte. 1,073,741,824 (2³⁰) bytes.

Gigabit Ethernet A group of Ethernet standards in which data is transmitted at 1 gigabit per second (Gbit/s). Abbreviated GbE.

gigabyte See **GB**.

group A collection of users who have similar needs. Groups simplify the administration of shared resources.

group folder A folder that organizes documents and applications of special interest to group members and allows group members to pass information among themselves.

guest computer A computer that doesn't have a computer account.

guest user A user who can log in to your server without a user name or password.

high availability The ability of a system to perform its function continuously, without interruption.

home directory See **home folder**.

home folder A folder for a user's personal use. Mac OS X also uses the home folder to store system preferences and managed user settings for Mac OS X users. Also known as a home directory.

host Another name for a server.

host name A unique name for a computer, historically referred to as the UNIX hostname.

HTML Hypertext Markup Language. The set of symbols or codes inserted in a file to be displayed on a web browser page. The markup tells the web browser how to display a webpage's words and images for the user.

HTTP Hypertext Transfer Protocol. The client/server protocol for the World Wide Web. HTTP provides a way for a web browser to access a web server and request hypermedia documents created using HTML.

Hypertext Markup Language See **HTML**.

Hypertext Transfer Protocol See **HTTP**.

IANA Internet Assigned Numbers Authority. An organization responsible for allocating IP addresses, assigning protocol parameters, and managing domain names.

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.

identity certificate See **certificate**.

IGMP Internet Group Management Protocol. An Internet protocol used by hosts and routers to send packets to lists of hosts that want to participate in a process known as multicasting. QuickTime Streaming Server (QTSS) uses multicast addressing, as does Service Location Protocol (SLP).

image See **disk image**.

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.

Internet A set of interconnected computer networks communicating through a common protocol (TCP/IP). The Internet is the most extensive publicly accessible system of interconnected computer networks in the world.

intranet A network of computers operated by and for the benefit of an organization's internal users. Access is commonly restricted to members of the organization. Many times, it refers to a website for the organization which is accessible only from within the organization. Intranets use the same networking technologies as the Internet (TCP/IP), and sometimes bridge legacy information systems with modern networking technologies.

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 data packets and TCP keeps track of data packets.

IP address A unique numeric address that identifies a computer on the Internet.

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.

IPP Internet Printing Protocol. A client-server protocol for printing over the Internet. The Mac OS X printing infrastructure and the Mac OS X Server print service that's built on it support IPP.

IPSec A security addition to IP. A protocol that provides data transmission security for L2TP VPN connections. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec nodes.

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.

journal data In Xsan, data about file system transactions that occur on an Xsan volume.

KB Kilobyte. 1,024 (2¹⁰) bytes.

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. After a user is authenticated, it's possible to access additional services without retyping a password (called single sign-on) for services that have been configured to take Kerberos tickets. Mac OS X Server uses Kerberos v5.

Kerberos Key Distribution Center See **KDC**.

Kerberos realm The authentication domain comprising the users and services that are registered with the same Kerberos server. The registered users and services trust the Kerberos server to verify each other's identities.

kilobyte See **KB**.

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.

layer A mechanism for prioritizing the tracks in a movie or the overlapping of sprites. When QuickTime plays a movie, it displays the movie's images according to their layer. Images with lower layer numbers are displayed on top; images with higher layer numbers may be obscured by images with lower layer numbers.

LDAP Lightweight Directory Access Protocol. A standard client-server protocol for accessing a directory domain.

lease period A limited period of time during which IP addresses are assigned. By using short leases, DHCP can reassign IP addresses on networks that have more computers than available IP addresses.

Lightweight Directory Access Protocol See **LDAP**.

link An active physical connection (electrical or optical) between two nodes on a network.

link aggregation Configuring several physical network links as a single logical link to improve the capacity and availability of network connections. With link aggregation, all ports are assigned the same ID. Compare to **multipathing**, in which each port keeps its own address.

load balancing The process of distributing client computers' requests for network services across multiple servers to optimize performance.

local area network See **LAN**.

local directory domain A directory of identification, authentication, authorization, and other administrative data that's accessible only on the computer where it resides. The local directory domain isn't accessible from other computers on the network.

local domain A directory domain that can be accessed only by the computer it resides on.

local home directory See **local home folder**.

local home folder A home folder that resides on disk on the computer a user is logged in to. It's accessible only by logging directly in to the computer where it resides, unless you log in to the computer using SSH.

local hostname A name that designates a computer on a local subnet. It can be used without a global DNS system to resolve names to IP addresses. It consists of lowercase letters, numbers, or hyphens (except as the last characters), and ends with ".local" (For example, bills-computer.local). Although the default name is derived from the computer name, a user can specify this name in the Sharing pane of System Preferences. It can be changed easily, and can be used anywhere a DNS name or fully qualified domain name is used. It can only resolve on the same subnet as the computer using it.

log in (verb) To start a session with a computer (often by authenticating as a user with an account on the computer) in order to obtain services or access files. Note that logging in is separate from connecting, which merely entails establishing a physical link with the computer.

long name The long form of a user or group name. See also **user name**.

LPR Line Printer Remote. A standard protocol for printing over TCP/IP.

MAC Media access control. See **MAC address**.

MAC address Media access control address. A hardware address that uniquely identifies each node on a network. For AirPort devices, the MAC address is called the AirPort ID.

Mac OS X The latest version of the Apple operating system. Mac OS X combines the reliability of UNIX with the ease of use of Macintosh.

Mac OS X Server An industrial-strength server platform that supports Mac, Windows, UNIX, and Linux clients out of the box and provides a suite of scalable workgroup and network services plus advanced remote management tools.

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.

managed preferences System or application preferences that are under administrative control. Workgroup Manager allows administrators to control settings for certain system preferences for Mac OS X managed clients.

master zone The DNS zone records held by a primary DNS server. A master zone is replicated by zone transfers to slave zones on secondary DNS servers.

MB Megabyte. 1,048,576 (2²⁰) bytes.

media access control See **MAC address**.

megabyte See **MB**.

migrate To transfer existing information, such as user and group accounts and user data, from one server or network to another server or network that's managed using different software.

mirrored Refers to a disk array that uses RAID 1, or mirroring.

mirroring Writing identical copies of data to two physical drives. Mirroring protects data against loss due to disk failure, and is the simplest method of achieving data redundancy.

mount (verb) To make a remote directory or volume available for access on a local system. In Xsan, to cause an Xsan volume to appear on a client's desktop, just like a local disk.

mount point In streaming, a string used to identify a live stream, which can be a relayed movie stream, a nonrelayed movie stream, or an MP3 stream. Mount points that describe live movie streams always end with a .sdp extension.

MS-CHAP Microsoft Challenge Handshake Authentication Protocol. The standard Windows authentication method for VPN. This authentication method encodes passwords when they are sent over the network and stores them in a scrambled form on the server. It offers good security during network transmission. MS-CHAP is a proprietary version of CHAP.

multicast DNS A protocol developed by Apple for automatic discovery of computers, devices, and services on IP networks. Called Bonjour (previously Rendezvous) by Apple, this proposed Internet standard protocol is sometimes referred to as ZeroConf or multicast DNS. For more information, visit www.apple.com or www.zeroconf.org. To see how this protocol is used in Mac OS X Server, see **local hostname**.

MySQL An open source relational database management tool frequently used by web servers.

name server A server on a network that keeps a list of names and the IP addresses associated with each name. See also **DNS**, **WINS**.

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.

network address translation See **NAT**.

Network File System See **NFS**.

Network Image Utility A utility provided with Mac OS X Server software that allows you to create disk images for NetBoot and Network Install services. Disk images can contain the Mac OS X operating system, applications, or both.

network installation The process of installing systems and software on Mac OS X client computers over the network. Software installation can occur with an administrator attending the installations or completely unattended.

network interface Your computer's hardware connection to a network. This includes (but isn't limited to) Ethernet connections, AirPort cards, and FireWire connections.

Network Time Protocol See **NTP**.

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 based on IP address, rather than user name and password.

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.

offline Refers to data that isn't immediately available, or to a device that is physically connected but not available for use.

online Refers to data, devices, or network connections that are available for immediate use.

Open Directory The Apple directory services architecture, which can access authoritative information about users and network resources from directory domains that use LDAP, Active Directory protocols, or 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 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.

package install image A file that you can use to install packages. Using NetBoot, client computers can start up over the network using this image to install software. Unlike block copy disk images, you can use same package install image for different hardware configurations.

partition A subdivision of the capacity of a physical or logical disk. Partitions are made up of contiguous blocks on the disk.

password An alphanumeric string used to authenticate the identity of a user or to authorize access to files or services.

password policy A set of rules that regulate the composition and validity of a user's password.

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 No Access. See also **privileges**.

PHP PHP Hypertext Preprocessor (originally Personal Home Page). A scripting language embedded in HTML that's used to create dynamic webpages.

plaintext Text that hasn't been encrypted.

Point to Point Tunneling Protocol See PPTP.

point-to-point One of three physical topologies that Fibre Channel uses to interconnect nodes. The point-to-point topology consists of a single connection between two nodes.

port A sort of virtual mail slot. A server uses port numbers to determine which application should receive data packets. Firewalls use port numbers to determine whether data packets are allowed to traverse a local network. "Port" usually refers to either a TCP or UDP port.

port name A unique identifier assigned to a Fibre Channel port.

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.

PPTP Point to Point Tunneling Protocol. A network transport protocol used for VPN connections. It's the Windows standard VPN protocol and uses the user-provided password to produce an encryption key.

private key One of two asymmetric keys used in a PKI security system. The private key is not distributed and is 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 and 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.

process A program that has started executing and has a portion of memory allocated to it.

protocol A set of rules that determines how data is sent back and forth between two applications.

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 cryptography A method of encrypting data that uses a pair of keys, one public and one private, that are obtained from a certification authority. One key is used to encrypt messages, and the other is used to decrypt them.

public key infrastructure A secure method of exchanging data over an unsecure public network, such as the Internet, by using public key cryptography.

QTSS Publisher An Apple application (included with Mac OS X Server) for managing QuickTime media and playlists, and preparing media for streaming and downloading.

QuickTime Streaming Server See **QTSS**.

RADIUS Remote Authentication Dial-In User Service.

RADIUS server A computer on the network that provides a centralized database of authentication information for computers on the network.

RAID Redundant Array of Independent (or Inexpensive) Disks. A grouping of multiple physical hard disks into a disk array, which either provides high-speed access to stored data, mirrors the data so that it can be rebuilt in case of disk failure, or both. The RAID array is presented to the storage system as a single logical storage unit. See also **RAID array**, **RAID level**.

RAID 0 A RAID scheme in which data is distributed evenly in stripes across an array of drives. RAID 0 increases the speed of data transfer, but provides no data protection.

RAID 0+1 A combination of RAID 0 and RAID 1. This RAID scheme is created by striping data across multiple pairs of mirrored drives.

RAID 1 A RAID scheme that creates a pair of mirrored drives with identical copies of the same data. It provides a high level of data availability.

RAID 5 A RAID scheme that distributes both data and parity information across an array of drives one block at a time, with each drive operating independently. This enables maximum read performance when accessing large files.

RAID array A group of physical disks organized and protected by a RAID scheme and presented by RAID hardware or software as a single logical disk. In Xsan, RAID arrays appear as LUNs, which are combined to form storage pools.

RAID set See **RAID array**.

realm General term with multiple applications. See **WebDAV realm**, **Kerberos realm**.

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.

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.

SACL Service Access Control List. Lets you specify which users and groups have access to specific services. See **ACL**.

Samba Open source software that provides file, print, authentication, authorization, name resolution, and network service browsing to Windows clients using the SMB protocol.

schema The collection of attributes and record types or classes that provide a blueprint for the information in a directory domain.

search base A distinguished name that identifies where to start searching for information in an LDAP directory's hierarchy of entries.

search path See **search policy**.

search policy A list of directory domains searched by a Mac OS X computer when it needs configuration information; also, the order in which domains are searched. Sometimes called a search path.

Secure Sockets Layer See **SSL**.

server A computer that provides services (such as file service, mail service, or web service) to other computers or network devices.

Server Message Block See **SMB**.

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.

shell A program that runs other programs. You can use a shell to interact with the computer by typing commands at a shell prompt. See also **command-line interface**.

short name An abbreviated name for a user. The short name is used by Mac OS X for home folders, authentication, and email addresses.

slave zone The DNS zone records held by a secondary DNS server. A slave zone receives its data by zone transfers from the master zone on the primary DNS server.

SLP DA Service Location Protocol Directory Agent. A protocol that registers services available on a network and gives users easy access to them. When a service is added to the network, the service uses SLP to register itself on the network. SLP DA uses a centralized repository for registered network services.

SMB Server Message Block. 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. SMB services use SMB 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 is usually 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.

stripe (noun) A partition of a drive in a RAID array.

stripe (verb) To write data to successive stripes in a RAID array or LUN.

subdirectory A directory within a directory.

subdomain Sometimes called the host name. Part of the domain name of a computer on the Internet. It does not include the domain or the top-level domain (TLD) designator (for example, .com, .net, .us, .uk). The domain name "www.example.com" consists of the subdomain "www," the domain "example," and the top-level domain "com."

subnet A grouping on the same network of client computers that are organized by location (for example, different floors of a building) or by usage (for example, all eighth-grade students). The use of subnets simplifies administration. See also **IP subnet**.

subnet mask A number used in IP networking to specify which portion of an IP address is the network number.

TB Terabyte. 1,099,511,627,776 (2⁴⁰) bytes.

TCP Transmission Control Protocol. A method used with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. IP handles the actual delivery of the data, and TCP keeps track of the units of data (called packets) into which a message is divided for efficient routing through the Internet.

terabyte See **TB**.

throughput The rate at which a computer can process data.

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.

URL Uniform Resource Locator. The address of a computer, file, or resource that can be accessed on a local network or the Internet. The URL is made up of the name of the protocol needed to access the resource, a domain name that identifies a specific computer on the Internet, and a hierarchical description of a file location on the computer.

user ID See **UID**.

user name The long name for a user, sometimes referred to as the user's real name. See also **short name**.

Virtual Private Network See **VPN**.

volume A mountable allocation of storage that behaves, from the client's perspective, like a local hard disk, hard disk partition, or network volume. In Xsan, a volume consists of one or more storage pools.

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 they 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 to the site while the site is running.

WebDAV realm A region of a website, usually a folder or directory, that's defined to provide access for WebDAV users and groups.

weblog See **blog**.

Weblog service The Mac OS X Server service that lets users and groups securely create and use blogs. Weblog service uses Open Directory authentication to verify the identity of blog authors and readers. If accessed using a website that's SSL enabled, Weblog service uses SSL encryption to further safeguard access to blogs.

wide area network See **WAN**.

wiki A website that allows users to collaboratively edit pages and easily access previous pages using a web browser.

Windows Internet Naming Service See **WINS**.

WINS Windows Internet Naming Service. A name resolution service used by Windows computers to match client names with IP addresses. A WINS server can be located on the local network or externally on the Internet.

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.

A

access

- ACLs 57, 73
- IP address restrictions 54
- Keychain Access Utility 66
- LDAP 21
- remote installation 84
- SACLs 73, 74
- user 145, 148
- See *also* permissions

accounts. See user accounts; Workgroup Manager

ACLs (access control lists) 57, 73

addresses. See IP addresses

Administer permission level 151

administrator 73, 74, 151

administrator computer 82, 138, 139

AFP (Apple Filing Protocol) service 22, 186

Apple Remote Desktop (ARD) 51, 144, 186

archiving server data 33, 36

ARD. See Apple Remote Desktop

`asx` tool 37, 87

authentication

- Kerberos 21, 59, 60, 112
- key-based SSH 71, 72
- keychain services 158
- MS-CHAPv2 110
- Open Directory 59
- overview 58
- passwords 59, 76, 98
- RADIUS 20, 22, 59, 157
- SASL 59
- Server Admin 40, 63, 140
- single sign-on 60
- standalone server 111
- and TLS 56
- users 58, 60, 72, 110
- Workgroup Manager 153
- See *also* certificates

authorization 58

See *also* authentication

B

backups

- advanced configuration 19
- command-line tools 37
- critical files 157
- media types 36
- policy considerations 32, 36
- rotation scheme 35
- scheduling 34
- server setup data 121
- types 33
- validation of 35

Berkeley Software Distribution. See BSD

broadcasting setup 135

BSD (Berkeley Software Distribution) 23

C

calendar service. See iCal service

Certificate Authority (CA)

- creating 65
- creating certificates from 67
- distributing to clients 69
- introduction 61
- overview 62
- requesting certificates from 63, 64, 65, 67
- See *also* PKI

Certificate Manager 62, 68

certificates

- creating 65, 67
- deleting 70
- editing 69
- identities 62
- importing 68
- managing 68
- overview 60, 61
- preparing 64
- private keys 61
- public keys 61
- renewing 70
- requesting 64
- root 65
- self-signed 62, 65, 69

- and Server Admin 62, 149
 - and services 70
 - Certificate Signing Request. *See* CSR
 - changeip tool 32
 - chat service. *See* iChat
 - client computers and NetBoot 28
 - clients
 - certificates 69
 - client-side logging 186
 - group accounts 154
 - and NetBoot 28
 - See also* users
 - command-line tools
 - backup tools 37
 - daemon control 171
 - disk space monitoring 175
 - erasing disks 97
 - installing server software 103
 - partitioning disks 94
 - and permissions 151
 - restoration tools 37
 - server administration 49
 - computer lists 153, 154
 - computer name 109, 144
 - computers, administrator 82, 138, 139
 - computer-to-computer network 166
 - computer-to-switch network 167
 - computer-to-switch-pair network 167
 - concatenated RAID set 95
 - configuration
 - advanced 19, 20, 112
 - authentication 59
 - automatic 117, 123, 124, 127
 - batch setup for multiple servers 115
 - connecting to network 108, 166, 167
 - DHCP 83, 112
 - directory connection 111, 112
 - Ethernet 108
 - interactive 112, 113, 114, 115
 - introduction 18, 107
 - link aggregation 168
 - logs 131
 - Open Directory 109, 110, 111, 112, 123, 127, 131
 - postponing 107
 - providing files to servers 122, 123
 - remote server 113, 114, 115
 - sample setup 187
 - saving setup data 118, 119, 120, 121, 124, 127
 - server infrastructure 30
 - server types 18
 - services 131, 132, 133, 134, 135, 136, 157
 - settings overview 109
 - SSL 149
 - standalone server 110
 - status checking 129, 130, 131
 - troubleshooting 130
 - types of 107, 146
 - worksheet for 197
 - Console 175
 - CSR (Certificate Signing Request) 63, 64, 65, 67
- ## D
- daemons, overview 171
 - Darwin (core operating system) 23
 - Date & Time preferences 145
 - d.f. tool 175
 - DHCP (Dynamic Host Configuration Protocol)
 - service 30, 83, 112
 - digital signature 149
 - directories. *See* directory services; domains, directory; folders
 - Directory, overview 44, 45
 - directory services
 - advanced configuration 112
 - and automatic setup 120, 123, 127
 - directory domains 21, 83, 110, 112, 156
 - logs 185
 - planning of 27
 - See also* Open Directory
 - Directory Utility 46
 - disk images
 - encrypting 58
 - installing with 28, 48, 87, 90
 - disks
 - command-line management of 94, 97, 175
 - erasing free space 97
 - installation preparation 91, 93, 94, 95, 96, 97
 - mirroring 95
 - monitoring tools 175
 - partitions 86, 93, 94, 95, 97
 - quotas 28
 - See also* RAID
 - diskspacemonitor tool 176
 - Disk Utility 58, 93, 95, 97
 - diskutil tool 94, 96, 97
 - ditto tool 37
 - DMZ, network 54
 - DNS (Domain Name System) service 30, 83
 - documentation 13, 14, 15
 - Domain Name System. *See* DNS
 - domains, directory 21, 83, 110, 112, 156
 - See also* Open Directory
 - drives. *See* disks
 - du tool 176
 - DVDs, installation 86
 - Dynamic Host Configuration Protocol. *See* DHCP
- ## E
- email. *See* mail service
 - emond daemon 183
 - encryption 56, 57, 61, 121

See also SSL
Ethereal packet sniffing tool 177
Ethernet 55, 108, 168
exporting service settings 148

F

files
 backup 32, 36, 157
 configuration 184
 full file-level copies 34
 security 57, 58
 setup data 118, 119, 120, 124
 shared secret 61
 storage considerations 28
file services 20, 22, 132, 186
file sharing 132, 150
file systems
 backing up 37
 choosing 91
 setup data 122
 See also volumes
File Transfer Protocol. *See* FTP
FileVault 57
Firewall service 54, 55, 83, 158
folders 27, 57, 144
FTP (File Transfer Protocol) service 22
full file-level copies 34
full image backup type 34

G

Gateway Setup Assistant 157
group accounts 154
groups 142, 148, 151, 153
Growl application 186

H

hardware requirements 17, 31, 81, 95
help, using 12
HFS+J volume 92
HFSX volume 92
historical data collection 173
home folders 27, 144
host name
 changing 146
 local 109, 144

I

iCal service 136, 158
iChat service 20, 136, 158
identity certificates. *See* certificates
images. *See* disk images; NetBoot; NetInstall
importing
 certificates 68
 service settings 148
incremental backups 34

infrastructure requirements 30
Inspector 156
installation
 administrator computer 82
 collecting information 81
 command-line method 103
 directory connections 83
 with disk images 28, 48, 87, 90
 disk preparation 91, 93, 94, 95, 96, 97
 from earlier OS versions 26, 28, 79, 82
 host name changing 146
 identifying servers 98
 infrastructure requirements 30
 integration strategy 29
 interactive 99, 101, 102
 multiple server 105
 network services setup 83
 overview 79
 planning for 25, 26, 27, 28, 29
 postponing setup after 107
 remote access 82, 84, 98, 101
 server installation disc 82
 server software 83, 103
 starting up for 83, 84, 86, 90
 system requirements 81
 updating 106
 upgrading 106

installer tool 103, 105

IP addresses

 access restriction 54
 changing server 32, 145
 and firewalls 83
 overview 23
 remote server installation 84, 98
 servers on different subnets 113
IPv6 addressing 23

J

journaling, file system 92

K

KDC (Kerberos Key Distribution Center). *See* Kerberos
Kerberos 21, 59, 60, 112
key-based authentication 71, 72
Keychain Access Utility 66
keychain services 158

L

LACP (Link Aggregation Control Protocol) 166
launchctl tool 172
launchd daemon 37, 171
LDAP (Lightweight Directory Access Protocol)
 service 21
LDAPv3 servers 59
link aggregation 165, 166, 167, 168, 169

- Link Aggregation Control Protocol. *See* LACP
- load balancing 170
- local directory domain, standalone server 111
- login, authenticating 71, 72
- logs
 - monitoring 175, 183, 184, 185, 186
 - troubleshooting setup 131
 - web services 161

M

- MAC (media access control) addresses 55, 98
- Mac OS X
 - administration from 139
 - installation considerations 82
 - upgrading from 106
- Mac OS X Server
 - administrative tools 39
 - configuration 110
 - integration strategy 29
 - introduction 17, 18
 - supported standards 21
 - system requirements 17
 - and UNIX 23
 - See also* configuration; installation
- mail service 20, 22, 134, 157, 159
- managed preferences, defining 154
- media, streaming. *See* streaming media
- migration 26, 28, 29
- mirroring, disk 95
- mobile accounts 144
- Monitor permission level 151
- MS-CHAPv2 authentication 110
- MySQL service 160

N

- Nagios application 186
- NAT (Network Address Translation) 159
- NetBoot service 28, 48, 90
- NetInstall 48, 90
- Network Address Translation. *See* NAT
- Network File System. *See* NFS
- network interfaces 144
- networks
 - connection configurations 108, 166, 167
 - environment for installation 80
 - Ethernet 55, 108, 168
 - initial server setup connection 108
 - monitoring tools 176, 179, 180, 181, 182
 - security 54, 55, 56
- network services
 - DHCP 30, 83, 112
 - DNS 30, 83
 - installation 83
 - NAT 159
 - NTP 144, 145

- planning for 30
- setup 134
- VLAN 55
- VPN 112
- See also* IP addresses
- network time protocol. *See* NTP
- NFS (Network File System) 22
- notification system 46, 145, 158, 177, 182
- See also* logs
- NTP (network time protocol) 144, 145

O

- Open Directory
 - authentication 59
 - logs 185
 - overview 20
 - and SACLs 73
 - setup 109, 110, 111, 112, 123, 127, 131
- Open Directory master 83
- Open Directory replica 59, 112, 164
- OpenLDAP 21
- open source modules
 - Kerberos 21, 59, 60, 112
 - OpenLDAP 21
 - OpenSSL 56
 - PHP 160
 - See also* Open Directory
- OpenSSL 56
- operating environment requirements 164

P

- PackageMaker 48
- packets, data, filtering of 54
- partitions, disk 86, 93, 94, 95, 97
- passwords 59, 76, 98
- permissions
 - administrator 73, 151
 - files 57
 - folder 57
 - SACL 74
 - types 57
- php configuration files 160
- physical infrastructure requirements 30
- PKI (public key infrastructure) 56, 60, 61
- Podcast Producer 135
- portable computers 144
- Portable Operating System Interface. *See* POSIX
- ports
 - Ethernet 108
 - list of 138
 - status of 138
 - TCP 70
- POSIX (Portable Operating System Interface) 57
- preferences 154
- presets 154

- print service 133
- private key 61, 62
- privileges, administrator 73, 151
 - See also permissions
- PropertyListEditor 48
- protocols
 - file service 22, 186
 - network service 30, 83, 112, 144, 145
 - overview 22
 - See also specific protocols
- public key certificates. See certificates
- public key cryptography 70
- public key infrastructure. See PKI

Q

- QuickTime Streaming Server (QTSS) 20, 49, 158
- quotas, disk space 28

R

- RADIUS (Remote Authentication Dial-In User Service) 20, 22, 59, 157
- RAID (Redundant Array of Independent Disks) 28, 94, 95, 96
- RAID Admin 175
- real-time monitoring 173
- Remote Authentication Dial-In User Service. See RADIUS
- remote servers
 - accessing 84
 - Apple Remote Desktop 51, 144, 186
 - configuration 113, 114, 115
 - identifying 98
 - installing from or to 82, 84, 98, 101
- replication 59, 112, 164
- requirements
 - hardware 17, 31, 81, 95
 - infrastructure 30
 - operating environment 164
 - software 81, 82
- restart, automatic 163
- restoration, data 32, 35
- root certificate 65
- rsync tool 37

S

- SACLs (service access control lists) 73, 74
- SASL (Simple Authentication and Security Layer) 59
- Secure Empty Trash 58
- secure SHell. See SSH
- Secure Sockets Layer. See SSL
- Secure VM 57
- security
 - administrator 73
 - authorization 58
 - best practices 74

- file 57, 58
- Firewall service 54, 55, 83, 158
- installation 83
- network 54, 55, 56
- overview 53
- physical 53
- SASL 59
- service level 73, 74
- settings 149
- SSH 70, 71, 72, 84, 85, 144, 158
- SSL 56, 60, 61, 62, 149
- TLS 56
- See also access; authentication; certificates; SSL
- self-signed certificates 62, 65, 69
- serial number, server 85
- Server Admin
 - access control 148
 - as administration tool 140, 141
 - authentication 40, 63, 140
 - certificates 62, 149
 - customizing 41
 - notification system 177
 - opening 40, 63, 140
 - overview 11, 39, 40, 63
 - server status 178
 - service management 147
 - and system imaging 48
- Server Assistant 42, 101, 107, 112
- Server Message Block protocol. See SMB
- Server Monitor 46, 174
- servers
 - adding 141
 - administration tools 39, 49, 50, 137, 140
 - basic settings 109, 143
 - groups of 142
 - infrastructure requirements 30
 - load balancing 170
 - reliability tools 161, 162, 163, 164, 165, 166, 168, 169
 - relocation considerations 31
 - removing 141
 - sample setup 187
 - serial numbers for 85
 - setup worksheet 197
 - standalone 109, 110, 111
 - startup 83, 90
 - status monitoring 173, 174, 175, 176, 177, 178
 - time 144, 145
 - troubleshooting 130
 - See also configuration; installation; remote servers
- Server Status Dashboard widget 174
- service access control lists. See SACLs
- services
 - access control 145, 148
 - exporting settings 148
 - importing settings 148

- management of 157
- planning for distribution of 27
- security 70, 73, 74
- setup 131, 132, 133, 134, 135, 136, 157
- viewing 145, 147
- See also* specific services
- setup procedures. *See* configuration; installation
- shared directory domain 21, 110
- shared secret files 61
- share points 57, 150
- Simple Network Management Protocol. *See* SNMP
- single points of failure 161
- single sign-on authentication 60
- slapd daemon 186
- SMB (Server Message Block) protocol 22
- snapshots, data 34
- SNMP (Simple Network Management Protocol)
 - definition 23
 - as monitoring tool 179, 180, 181, 182
 - settings 144
- snmpd daemon 180
- Software Update service 106, 135
- srm UNIX utility 58
- SSH (secure SHell host) 70, 71, 72, 84, 85, 144, 158
- SSL (Secure Sockets Layer) 56, 60, 61, 62, 149
- standalone server 109, 110, 111
- standard configuration type 18
- streaming media 20, 28, 49, 135, 158
- striping 95
- subnets 108, 113
- syslog configuration file 184
- syslogd daemon 183
- System Image Utility 48
- system imaging service 135

T

- TCP (Transmission Control Protocol) 54, 70
- `tcpdump` tool 177
- time server 144, 145
- TLS (Transport Layer Security) protocol 56
- Transmission Control Protocol. *See* TCP
- Transport Layer Security protocol. *See* TLS
- troubleshooting server operation 130

U

- UDP (User Datagram Protocol) 54
- UNIX 23
- upgrading
 - from Mac OS X 106
 - from previous server versions 26, 28
 - vs. migration 26, 29
 - and saved setup data 118
- UPS (uninterruptible power supply) 162, 163
- user accounts

- authentication 60
- group 154
- managed preferences 154
- management of 153
- mobile 144
- passwords 59
- setup 132
- See also* users
- User Datagram Protocol. *See* UDP
- users
 - access control 145, 148
 - administrative access for 73
 - authentication 58, 60, 72, 110
 - certificates 62
 - and Directory 44
 - disk space quotas 28
 - groups 148, 151, 153
 - home folders 27, 144
 - management of 153
 - permissions 151
 - Windows 28
 - See also* clients; user accounts; Workgroup Manager

V

- Virtual Private Network. *See* VPN
- VLAN (virtual local area network) 55
- VNC (virtual network computing) 81, 84, 102, 105
- volumes
 - backing up 37
 - erasing 97
 - and partitioning 93, 94
 - RAID 95
 - setup data 122
 - startup 84, 90
 - supported 92
- VPN (Virtual Private Network) 112

W

- weblog service 161
- WebObjects Application Server 136
- web services 20, 21, 133, 160
- web technologies 22
- wikis 161
- Windows NT 29
- Windows users 28
- workgroup configuration type 18
- Workgroup Manager
 - administering accounts 153
 - administration overview 152
 - authentication 153
 - customizing 44, 156
 - opening 42, 153
 - overview 42, 43

X

Xgrid 2 service 20, 157
Xgrid Admin 50
Xsan 19
Xserve

hardware installation instructions 81
and Server Monitor 46
and server reliability 162, 163
VLAN support 55