



Mac OS X Server Advanced Server Administration

For Version 10.6 Snow Leopard

© 2009 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, iDVD, iMovie, iPhoto, iPod, iTunes, Mac, Macintosh, the Mac logo, Mac OS, PowerBook, QuickTime and SuperDrive are trademarks of Apple Inc., registered in the U.S. and other countries.

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.

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.

019-1410/2009-05-29

Contents

11	Preface: About This Guide
11	What's in This Guide
12	Using Onscreen Help
12	Document Road Map
13	Viewing PDF Guides Onscreen
14	Printing PDF Guides
14	Getting Documentation Updates
14	Getting Additional Information
16	Chapter 1: System Overview and Supported Standards
16	System Requirements for Installing Mac OS X Server
16	What's New in Snow Leopard Server
18	What's New in Server Admin
18	Understanding Server Configuration Methods
19	Mac OS X Snow Leopard Server Enhancements
20	Supported Standards
23	Mac OS X Server's UNIX Heritage
24	Chapter 2: Planning
24	Planning
25	Planning for Upgrading or Migrating to Mac OS X Server v10.6
25	Setting Up a Planning Team
26	Identifying the Servers You'll Need to Set Up
26	Determining Services to Host on Each Server
27	Defining a Migration Strategy
28	Upgrading and Migrating from an Earlier Version of Mac OS X Server
28	Migrating from Windows
28	Defining an Integration Strategy
29	Defining Physical Infrastructure Requirements
29	Defining Server Setup Infrastructure Requirements
30	Making Sure Required Server Hardware Is Available
31	Minimizing the Need to Relocate Servers After Setup
31	Defining Backup and Restore Policies

31	Understanding Backup and Restore Policies
33	Understanding Backup Types
33	Understanding Backup Scheduling
34	Understanding Restores
35	Other Backup Policy Considerations
36	Command-Line Backup and Restoration Tools
36	Understanding Time Machine as a Backup Tool

38 Chapter 3: Administration Tools

38	Server Admin
38	Opening and Authenticating in Server Admin
39	Server Admin Interface
40	Customizing the Server Admin Environment
41	Server Assistant
41	Server Preferences
42	Workgroup Manager
43	Workgroup Manager Interface
44	Customizing the Workgroup Manager Environment
44	Server Monitor
45	iCal Service Utility
46	iCal Service Utility Interface
47	System Image Management
47	Media Streaming Management
48	Command-Line Tools
48	Server Status Widget
48	RAID Admin
49	Podcast Composer
49	Xgrid Admin
49	Apple Remote Desktop

51 Chapter 4: Security

51	About Physical Security
52	About Network Security
52	Firewalls and Packet Filters
52	Network DMZ
53	VLANs
53	MAC Filtering
54	Transport Encryption
55	Payload Encryption
55	About File Security
55	File and Folder Permissions
56	About File Encryption
56	Secure Delete

56	About Authentication and Authorization
58	Single Sign-On
59	About Certificates, SSL, and Public Key Infrastructure
60	Public and Private Keys
60	Certificates
61	About Certificate Authorities (CAs)
61	About Identities
61	About Self-Signed Certificates
61	About Intermediate Trust
63	Certificate Manager in Server Admin
64	Readying Certificates
65	Creating a Self-Signed Certificate
65	Requesting a Certificate From a Certificate Authority
66	Creating a Certificate Authority
68	Using a CA to Create a Certificate for Someone Else
68	Importing a Certificate Identity
69	Managing Certificates
70	Editing a Certificate
70	Distributing a CA Public Certificate to Clients
71	Deleting a Certificate
71	Renewing an Expiring Certificate
71	Replacing an Existing Certificate
72	Using Certificates
72	SSH and SSH Keys
72	Key-Based SSH Login
72	Generating a Key Pair for SSH
74	Administration Level Security
75	Setting Administration Level Privileges
75	Service Level Security
75	Setting SAFL Permissions
76	Security Best Practices
78	Password Guidelines
78	Creating Complex Passwords
79	Chapter 5: Installation and Deployment
79	Installation Overview
80	System Requirements for Installing Mac OS X Server
81	Hardware-Specific Instructions for Installing Mac OS X Server
81	Gathering the Information You Need
81	Setting Up Network Services
81	Connecting to the Directory During Installation
81	SSH Service During Installation
82	About The Server Installation Disc

82	Preparing an Administrator Computer
83	About Starting Up for Installation
83	Before Starting Up
83	Remotely Accessing the Install DVD
85	About Server Serial Numbers for Default Installation Passwords
85	Starting Up from the Install DVD
86	Starting Up from an Alternate Partition
89	Starting Up from a NetBoot Environment
90	Preparing Disks for Installing Mac OS X Server
90	Choosing a File System
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 Screen Sharing and VNC
103	Changing a Remote Computer's Startup Disk
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
107	How to Keep Current
108	Chapter 6: Initial Server Setup
108	Information You Need
108	Postponing Server Setup Following Installation
109	Connecting to the Network During Initial Server Setup
109	Configuring Servers with Multiple Ethernet Ports
109	About Settings Established During Initial Server Setup
110	Specifying Initial Open Directory Usage
111	Not Changing Directory Usage When Upgrading
112	Setting Up a Server as a Standalone Server
112	Setting Up a Server to Connect to a Directory System
113	Binding a Server to Multiple Directory Servers
113	Using Interactive Server Setup
114	Setting Up a Local Server Interactively
115	Setting Up a Remote Server Interactively
116	Setting Up Multiple Remote Servers Interactively in a Batch
118	Using Automatic Server Setup
119	How a Server Searches for Saved Setup Data
120	Creating and Saving Setup Data
121	Using Encryption with Setup Data Files
121	Setting Up Servers Automatically Using Data Saved in a File
122	Setting a Mac OS X Server Serial Number with the Command-line
122	Determining the Status of Setups

122	Using the Destination Pane for Setup Status Information
123	Handling Setup Failures
123	Handling Setup Warnings
123	Getting Upgrade Installation Status Information
124	Setting Up Services
124	Adding Services to the Server View
124	Setting Up Open Directory
125	Setting Up User Management
125	Setting Up All Other Services
126	Chapter 7: Management
126	Computers You Can Use to Administer a Server
126	Setting Up an Administrator Computer
127	Using a Non-Mac OS X Computer for Administration
127	Using the Administration Tools
128	Working With Pre-Version 10.6 Computers From Version 10.6 Servers
129	Ports Used for Administration
129	Ports Open By Default
129	Server Admin Basics
130	Adding and Removing Servers in Server Admin
131	Grouping Servers Manually
131	Grouping Servers Using Smart Groups
132	Working with Settings for a Specific Server
134	Understanding Changes to the Server IP Address or Network Identity
135	Understanding Mac OS X Server Names
135	Understanding IP Address or Network Identity Changes on Infrastructure Services
138	Understanding IP Address or Network Identity Changes on Web and Wiki Services
139	Understanding IP Address or Network Identity Changes on File Services
140	Understanding IP Address or Network Identity Changes on Mail Services
142	Understanding IP Address or Network Identity Changes on Collaboration Services
143	Understanding IP Address or Network Identity Changes on Podcast Producer
145	Understanding IP Address or Network Identity Changes on Other Services
146	Changing the IP Address of a Server
146	Changing the Server's DNS Name After Setup
146	Changing the Server's Computer Name and the Local Hostname
147	Administering Services
148	Adding and Removing Services in Server Admin
149	Importing and Exporting Service Settings
149	Controlling Access to Services
150	Using SSL for Remote Server Administration
150	Managing Sharing
151	Tiered Administration Permissions
152	Defining Administrative Permissions

152	Workgroup Manager Basics
153	Opening and Authenticating in Workgroup Manager
153	Administering Accounts
154	Working with Users and Groups
155	Defining Managed Preferences
155	Working with Directory Data
156	Customizing the Workgroup Manager Environment
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
164	Link Aggregation
165	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
173	Chapter 8: Monitoring
173	Planning a Monitoring Policy
174	Planning Monitoring Response
174	Server Status Widget
174	Server Monitor
175	RAID Admin
175	Console
175	Disk Monitoring Tools
177	Network Monitoring Tools
177	Server Status Notification in Server Admin
178	Monitoring Server Status Overviews Using Server Admin
178	Remote Kernel Core Dumps
180	Setting Up a Core Dump Server
181	Setting Up a Core Dump Client
182	Configuring Common Core Dump Options
182	Simple Network Management Protocol (SNMP)
183	Enabling SNMP reporting (TBD update)

183	Configuring snmpd
185	Additional Information about SNMP
185	Tools to Use with SNMP
186	Notification and Event Monitoring Daemons
187	Logging
187	Syslog
188	Directory Service Debug Logging
189	Open Directory Logging
189	AFP Logging
189	Additional Monitoring Aids
191	Chapter 9: Push Notification Server
191	Provide increased server responsiveness to clients and reduce server load with Push Notification Server.
191	Starting and Stopping Push Notification
193	Index

About This Guide

This guide provides a starting point for administering Mac OS X Leopard Server using its advanced administration tools. 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 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” gives you advice for planning Mac OS X Server deployment.
- 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 Snow Leopard Server. You can view help on a server or an administrator computer. (An administrator computer is a Mac OS X computer with Snow Leopard Server administration software installed on it.)

To get the most recent onscreen help for Mac OS X Snow 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 later.

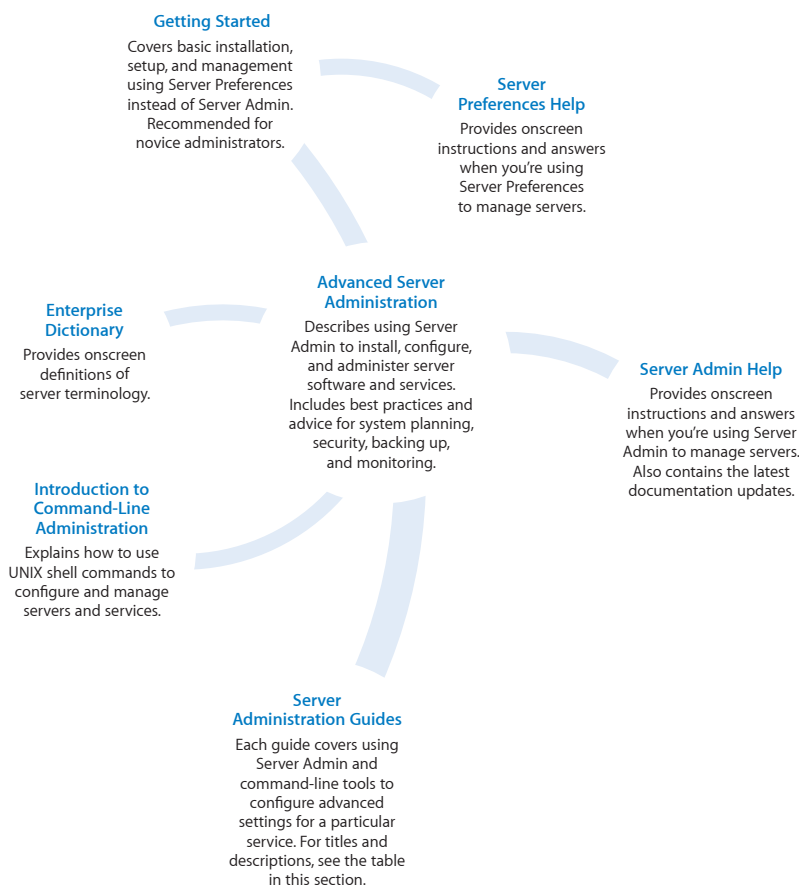
To see the most recent server help topics:

- 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.

Document Road Map

Snow Leopard has a suite of guides which can cover management of individual services. Each service may be dependent on other services for maximum utility. The road map below shows some related documentation that you may need to fully configure your desired service to your specifications. You can get these guides in PDF format from the Mac OS X Server documentation website:



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/resources/
- 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/snowleopard/serverdocupdates.xml

Getting Additional Information

For more information, consult these resources:

- *Read Me documents*—get important updates and special information. Look for them on the server discs.
- *Mac OS X Server website* (www.apple.com/server/macosx/)—enter the gateway to extensive product and technology information.

- *Mac OS X Server Support website* (www.apple.com/support/macosxserver/)—access hundreds of articles from Apple's support organization.
- *Apple Discussions website* (discussions.apple.com/)—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.
- *Apple Training and Certification website* (www.apple.com/training/)—hone your server administration skills with instructor-led or self-paced training, and differentiate yourself with certification.

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 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 v10.6 Snow Leopard must have:

- An Intel processor, 1.5 GHz or faster
- 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 Server 10.6

The serial number used with any previous version of Mac OS X Server will not allow registration in v10.6.

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 126.

What's New in Snow Leopard Server

Snow Leopard Server offers major enhancements in several key areas:

- Address Book Server

Snow Leopard Server introduces the first open standards-based Address Book Server Based on the emerging CardDAV specification, which uses WebDAV to exchange vCards, sharing contacts across multiple computers.

- Remote Access
Snow Leopard Server delivers push notifications to mobile users outside your firewall, and a proxy service gives them secure remote access to email, address book contacts, calendars, and select internal websites.
- Collaboration services improvements
Snow Leopard Server augments collaboration features with wiki and blog templates optimized for viewing on iPhone; content searching across multiple wikis; and attachment viewing in Quick Look. It also introduces My Page, which gives users one convenient place to access their web applications, receive notifications, and view activity streams across wikis.
- iCal Server 2
Snow Leopard Server's new iCal Server includes group and shared calendars, push notifications, the ability to send email invitations to non-iCal Server users, and a browser-based application for using calendars with many supported browsers.
- Podcast Producer 2
Snow Leopard Server's new Podcast Producer features an intuitive new workflow editor, support for dual-video source capture, and now includes Podcast Library, which lets you host locally stored podcasts and make them available for subscription by category via automatically generated Atom web feeds.
- Mail Server improvements
Snow Leopard Server mail service increases its performance and scalability using a new engine designed to handle thousands of simultaneous connections. Mail services have been enhanced to include server-side email rules and vacation messages.
- Multicore optimizations
Snow Leopard Server supports "Grand Central," a new set of built-in technologies that makes all of Mac OS X Server multicore aware and optimizes it for allocating tasks across multiple cores and processors.
- 64-bit support
Snow Leopard Server use 64-bit kernel technology to support up to 16TB of memory.
- OpenCL support
Snow Leopard Server supports OpenCL (Open Computing Language), makes it possible for developers to use the graphics processing unit (GPU) for general computational tasks.

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, streamlined, and integrated Server Assistant
- Smoother interaction with Server Preferences settings
- Improved user interface

Understanding Server Configuration Methods

Mac OS X Server can configure and manage Mac OS X Server using two configuration methods: Server Preferences, or the advanced configuration tool suite, which includes Server Admin. Servers administered using the advanced tool suite are the most flexible and require the most skill to administer. Servers administered by Server Preferences have fewer configuration options, but most of the configuration details are set by the Server Preferences, without additional skill or labor. You can customize your server for a variety of purposes using either method.

Using Server Admin and the rest of the advanced configuration tool suite, the experienced system administrator has complete control of each service's configuration to accommodate a wide variety of 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.

Using Server Preferences, you can get standard configurations of Mac OS X Server features using automated setup and simplified administration. For more information about using Server Preferences to administer your server, see *Getting Started*.

You may switch between Server Admin and Server Preferences as you see fit. The setting changes in one application are reflected in the other's settings. However, some advanced or custom configurations can't be inspected or changed in Server Preferences, due to Server Preferences' simplified interface.

The following table highlights the capabilities of each configuration method.

	Advanced tool suite	Server Preferences
Service settings changed with...	Server Admin	Server Preferences
Initial service settings are...	Unconfigured	Preset to common defaults
Users and groups managed with...	Workgroup Manager	Server Preferences

	Advanced tool suite	Server Preferences
User service access control managed with...	Server Admin	Server Preferences
Xserve diagnostics monitored with...	Server Monitor	Server Monitor
Xgrid clusters administered with...	Xgrid Admin	Xgrid Admin
Staus monitored and and settings backed up using...	Command line tools and third-party backup solutions implemented by the system administrator	Server Preferences
Your data (websites, databases, calendar files, etc.) backed up using...	Command line tools and third-party backup solutions implemented by the system administrator	Command line tools and third-party backup solutions implemented by the system administrator

Mac OS X Snow Leopard Server Enhancements

Server Admin offers major enhancements in several key areas:

- **Address Book Server**
Snow Leopard Server introduces the first open standards-based Address Book Server Based on the emerging CardDAV specification, which uses WebDAV to exchange vCards, sharing contacts across multiple computers.
- **Mobile Access**
Snow Leopard Server delivers push notifications to users outside your firewall, and a proxy service gives them secure remote access to email, address book contacts, calendars, and select internal websites.
- **Collaboration services improvements**
Snow Leopard Server augments collaboration features with wiki and blog templates optimized for viewing on iPhone; content searching across multiple wikis; and attachment viewing in Quick Look. It also introduces My Page, which gives users one convenient place to access their web applications, receive notifications, and view activity streams across wikis.
- **iCal Server 2**
Snow Leopard Server's next major release of iCal Server includes group and shared calendars, push notifications, the ability to send email invitations to non-iCal Server users, and a browser-based application for using calendars with many supported browsers.
- **Podcast Producer 2**

Snow Leopard Server's next major release of Podcast Producer features an intuitive new workflow editor, support for dual-video source capture, and now includes Podcast Library, which lets you host locally stored podcasts and make them available for subscription by category via automatically generated Atom web feeds.

- Mail Server improvements

Snow Leopard Server mail service increases its performance and scalability using a new engine designed to handle thousands of simultaneous connections. Mail services have been enhanced to include server-side email rules and vacation messages.

- Multicore optimizations

Snow Leopard Server supports “Grand Central,” a new set of built-in technologies that makes all of Mac OS X Server multicore aware and optimizes it for allocating tasks across multiple cores and processors.

- 64-bit support

Snow Leopard Server use 64-bit kernel technology to support up to 16TB of memory.

- OpenCL support

Snow Leopard supports OpenCL (Open Computing Language), makes it possible for developers to use the graphics processing unit (GPU) for general computational tasks.

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, Dovecot, 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 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 is a complete AMP stack (a bundle of integrated Apache-MySQL-PHP/Perl/Python software). 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.

- **XMPP:** Extensible Messaging and Presence Protocol (XMPP) is an open XML-based messaging protocol used for messaging and presence information. XMPP serves as the basis for Mac OS X Server's Push Notification service, as well as iChat Server, and all publish and subscribe functions in the server.

Mac OS X Server's UNIX Heritage

Mac OS X Server has a UNIX 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.

Planning

2

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 v10.6 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 using Server Preferences.

Planning for Upgrading or Migrating to Mac OS X Server v10.6

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 v10.6.

You can upgrade to Snow Leopard Server if you're using the latest update of Mac OS X Server v10.5 Leopard or Mac OS X Server v10.4.11 Tiger on a Mac OS X Servers with Intel processors, 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 Snow Leopard Server on.

Migration is supported from the latest update of Mac OS X Server v10.5 Leopard or Mac OS X Server v10.4.11 Tiger. 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 *Installation and Setup Worksheet*.

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 10.6?
- 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 to the clients and can be provided by a different server than the NetBoot server. 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 v10.4–10.5 or a Windows NT server, examine the opportunities for moving data and settings to Mac OS X Server v10.6.

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

If you're using computers with Mac OS X Server v10.4 or v10.5, consider upgrading or migrating them to Mac OS X Server v10.6 Snow Leopard.

If you're using Mac OS X Server v10.5 or v10.4 and you don't need to move to Intel-processor based 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 10.4 or 10.5 server's hard disk needs reformatting or the server doesn't meet the minimum Snow 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 v10.4 or 10.5 server to different server hardware.

Migration is supported from the latest versions of Mac OS X Server v10.5 and v10.4. When you migrate, you install and set up Snow Leopard Server, then restore files onto it from the earlier server, and make manual adjustments as required.

For complete information, read *Upgrading and Migrating*.

Migrating from Windows

Snow 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, Snow Leopard Server can replace Windows 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 non-Apple 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 physical 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?
- Are you considering moving your servers to different IP addresses or hostnames? Now may be a good time to do it.

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 in a directory 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:

Setting up basic server infrastructure:

- 1 Set up the DNS server.

Populate the DNS with the host names of the desired servers and services.

- 2 Set up DHCP.

Configure DHCP to specify the DNS server address so it can be served to DHCP clients. If desired, set up DHCP-managed static IP address for the servers.

- 3 Set up a directory server, including Windows PDC service if required.

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.

- 4 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 DNS 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).

You probably will need to manually adjust service and system settings. For more information on how to do this, see “Understanding Changes to the Server IP Address or Network Identity” on page 134.

- 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 can 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?
- What should not be backed up (as per organization policy)?
- 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, which include, but are not limited to:

- `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. By default `rsync` does not preserve extended attributes in files necessary for many Mac OS X Server services.
- `ditto`. Use this command to perform full backups.
- `asr`. Use this command to back up and restore an entire volume in block copy mode. If the tool is in file copy mode, it does not preserve necessary extended attributes in files.

For more information about these commands, see their respective man pages.

Note: You can use the `launchctl` command to automate data backup using the aforementioned commands. For more information about using `launchctl` and `launchd`, see their respective man pages.

Understanding Time Machine as a Backup Tool

At its core, Time Machine is a file-level backup solution which runs at regular intervals and archives file changes from the initial file set. Time Machine makes use of UNIX file linking to efficiently store backup intervals as separate browsable file systems, but uses no compression.

Time Machine is a limited tool for data backup and restoration of Mac OS X Server v10.6.

Time Machine can be used to backup certain server configuration settings and the service state. Time Machine does not back up service data. For example, Time Machine doesn't back up user and group directory records, email, DNS records, Address Book shared groups, iCal Server calendars, and so forth. It only saves the settings made in Server Preferences and Server Admin, and whether or not a service is on or off. The following service settings and statuses are preserved:

- Address Book Server
- DHCP
- DNS
- File Services (AFP, SMB, NFS, and FTP)
- Firewall
- iCal Server

- iChat Server
- Mail
- Mobile Access
- MySQL
- NAT
- Network Settings
- Podcast Producer
- Print
- Push Notification
- QTSS
- RADIUS
- Remote Access Settings
- Software Update
- VPN
- Web
- Wiki
- Xgrid

For more information about where the necessary data files are stored for backup via other means, see .

Note: You can use the `launchctl` command to automate data backup using the aforementioned commands. For more information about using `launchctl` and `launchd`, see their respective man pages.

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

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.

The version of Server Admin included with Mac OS X Server v10.6 can be used to administer the latest version of Mac OS X Server v10.5. However, the current version of Server Admin isn't compatible with administering DNS service or manage certificates in Mac OS X Server v10.5; use the version of Server Admin that came with Mac OS X Server v10.5 on a computer running Mac OS X Server v10.5 or Mac OS X v10.5.

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.

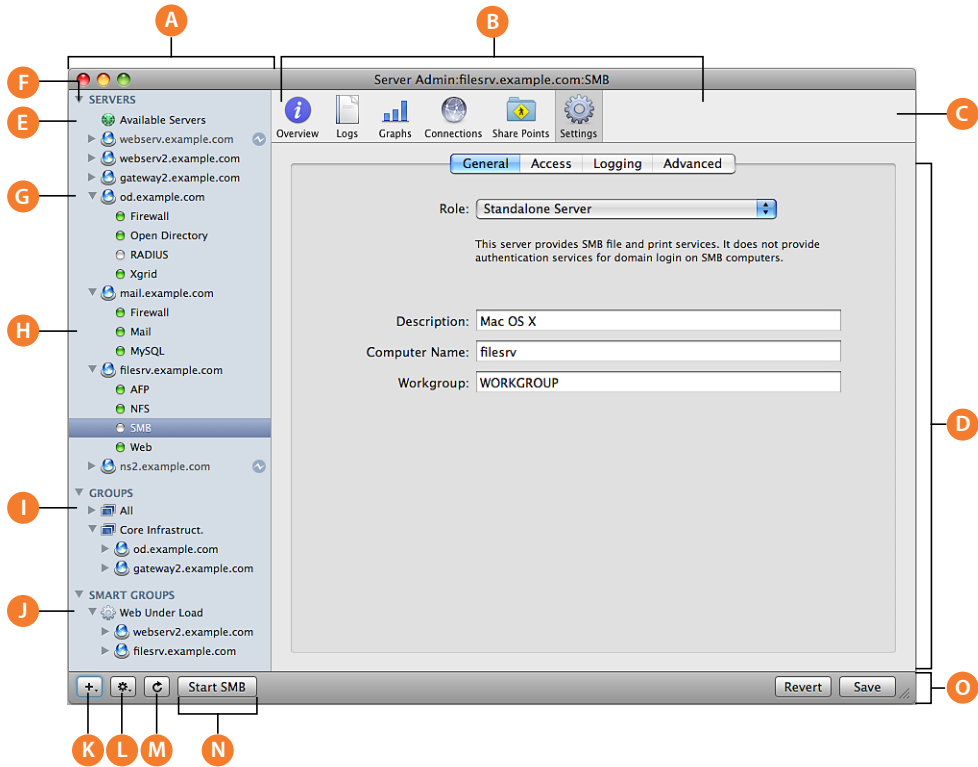
Opening and Authenticating in Server Admin

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 131.
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 131.
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.

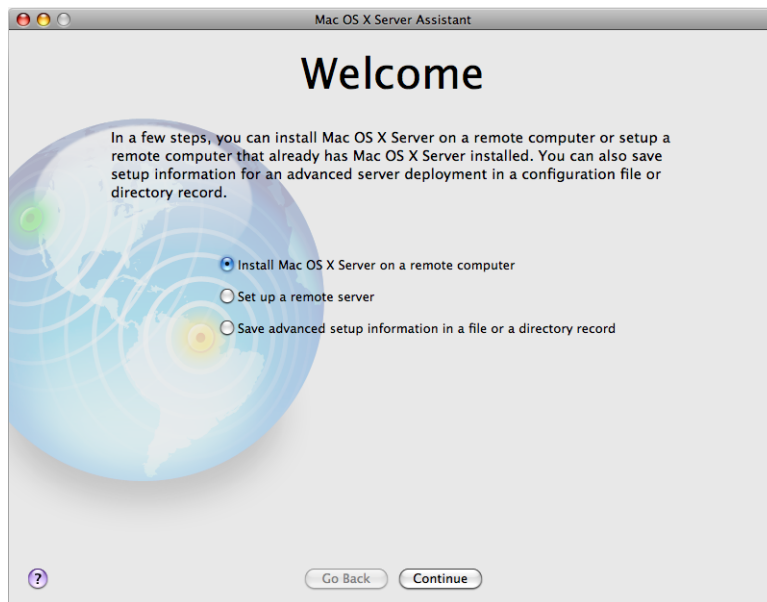
- To control the list of services to administer, see “Adding and Removing Servers in Server Admin” on page 130.
- To control the appearance of Server Admin lists, refresh rates, and other behaviors, choose Server Admin > Preferences.
- To group and sort servers available for administration, make groups and smart groups. See “Grouping Servers Manually” and “Grouping Servers Using Smart Groups.”

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

The Server Assistant initial page is shown here.



Server Assistant is launched from the Server menu of Server Admin. The following menu items launch the assistant:

- Install Remote Servers
- Configure Remote Servers
- Create Auto Server Setup Profile

For information about using Server Assistant, use its Help buttons, or see Chapter 6, "Initial Server Setup," on page 108.

Server Preferences

Server Preferences is the simplified administration application you need for managing Snow Leopard Server. Use Server Preferences instead of Server Admin to:

- Manage users and groups.

- Configure file sharing service, iCal calendar service, iChat instant messaging service, mail service, web services, VPN remote access service, and Time Machine backup for users' computers.
- Check the status of the server and services.

You can use Server Preferences on any server you want to manage, or you can use it remotely from an administrator computer or another server.

For information about using Server Preferences, see *Getting Started* or *Server Preferences Help*.

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:

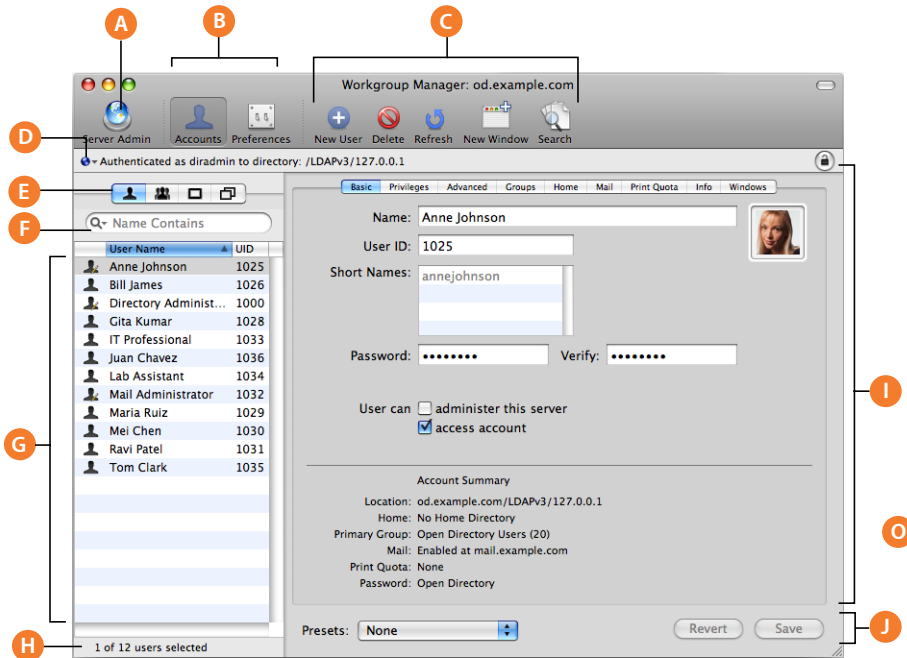
- *User Management* explains how to use Workgroup Manager for account and preference management. This guide also explains how to import and export accounts.
- *Open Directory Administration* describes how to use the Inspector.

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.

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 computer to monitor, click Add Server, identify the server, and enter user name and password information for an administrator of the server. If adding the local server, use '127.0.0.1' for the IP address. If adding a remote server, enter the server's LOM hostname or IP address.

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 computers 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 computer 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 computer 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 computer'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 computer. 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.)

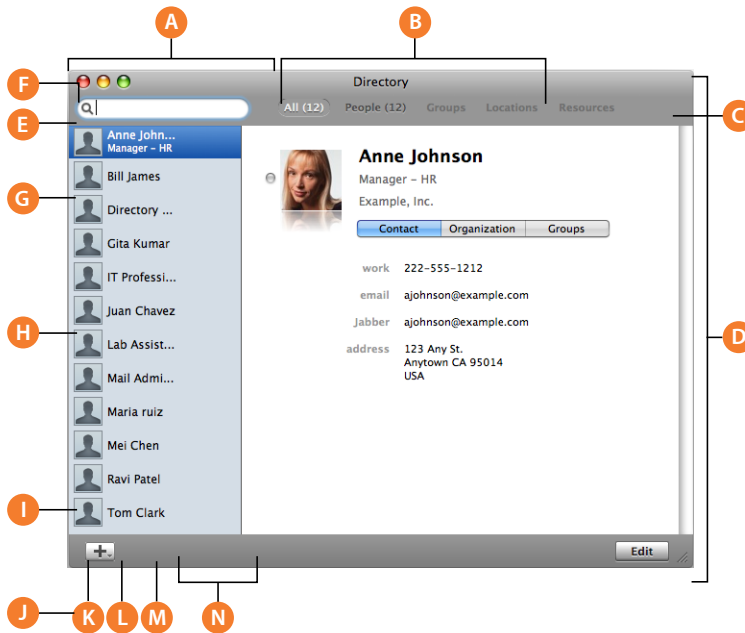
For additional information, see Server Monitor Help.

iCal Service Utility

iCal Service Utility gives users access to shared information about locations, and resources within the organization. They can use iCal Service Utility to set up information about shared resources and locations for use with iCal Service.

iCal Service Utility Interface

The iCal Service Utility interface is shown here, with each element explained in the following table.



NEED TO ADD ORANGE CIRCLES AND LETTERS



- 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 location or resource record.
- F** **Edit button:** Click to edit the selected record.

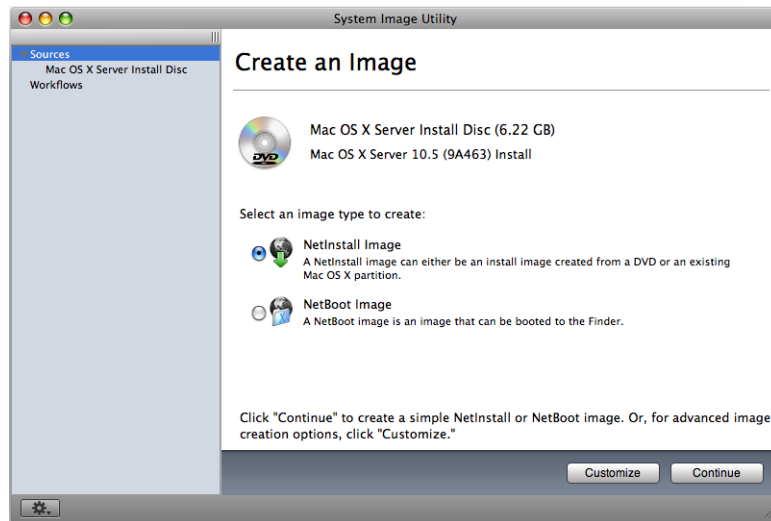
For information about how to use iCal Service Utility, see the onscreen help for iCal Service Utility.

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 Broadcaster, an easy-to-use application for producing live streaming QuickTime events.

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.

Mac OS X Server has a command-line version of Server Admin called `serveradmin` which can be used to administer all the services that Server Admin manages. It is run on the server to be administered over a remote connection.

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

Command-Line Administration describes how you can learn to use command line tools write scripts which automate and improve your processes.

Server Status Widget

The Server Status widget lets you monitor Snow Leopard Server activity from any computer with Snow Leopard or Snow Leopard Server. Server Status shows you graphs of processor activity, network load, and disk usage.

For information about using the Server Status widget, see *Getting Started* or Server Preferences Help.

RAID Admin

RAID Admin is a tool to administer and monitor Xserve RAID devices. You use RAID Admin to set up the Xserve RAID hardware, including:

- Creating, deleting, and expanding RAID arrays
- Monitoring the status of one or more Xserve RAID systems
- Adjusting settings, including system name and password, network address for each RAID controller, fibre channel communication speed, drive cache, and controller cache
- Setting up email notification for system alerts
- Implementing advanced features, such as dividing arrays into “slices” and updating the firmware of an Xserve RAID system.

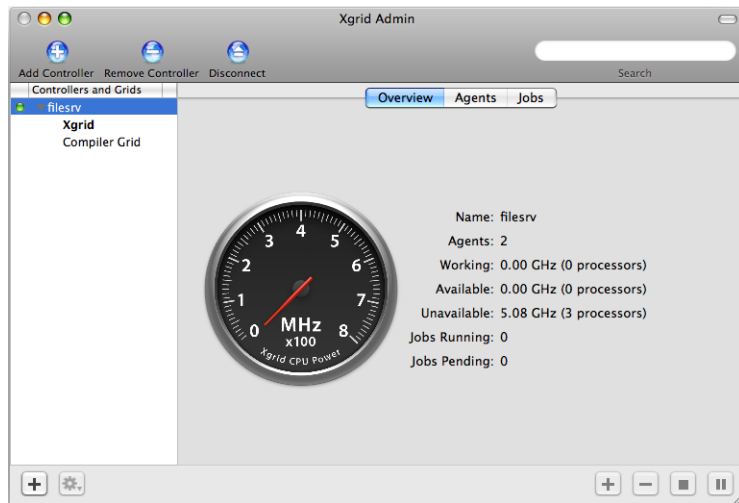
Podcast Composer

Podcast Composer records high-quality audio and video from a local or remote camera, capture screen activity, or upload QuickTime files into Podcast Composer for encoding and distribution.

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.



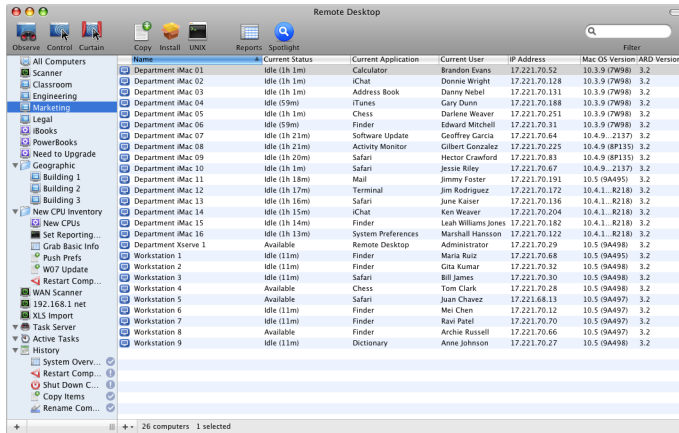
Xgrid Admin is installed in `/Applications/Server/` when you install your server or set up an administrator computer. To open Xgrid Admin, double-click the Xgrid Admin icon in `/Applications/Server/`.

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 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 42-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:

- Server administration using Server Admin and Server Preferences
- User and group management using Workgroup Manager.
- Address Book Server
- iCal Server
- iChat Server
- Mail Service
- Open Directory
- Podcast Producer
- RADIUS
- SSH
- VPN (L2TP)
- Web Service

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 56.

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 Else.

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 is allowed 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:

- Address Book Server
- Apple Filing Protocol (AFP)
- File Transfer Protocol (FTP)

- iCal Server
- iChat Server
- Login window
- Mail Services
- Network Filing Protocol (NFS)
- Open Directory (LDAPv3)
- Printing (IPP)
- Screen saver
- Secure Shell (SSH)
- Server Message Block file service (SMB)
- Virtual Private Network (VPN)
- Virtual Network Computing (VNC, known as Screen Sharing in Mac OS X Server)
- Web Service (Apache via the SPNEGO Simple and Protected GSS-API Negotiation Mechanism protocol)
- Xgrid
- **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 10.6 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

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

A certificate is an electronic document which contains a the public key with identification information (name, organization, email address, and so on). In a public key environment, a certificate is digitally “signed” by a Certificate Authority, or its own private key (the latter being a “self-signed certificate”).

A public key 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

About 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..

About Identities

Identities are a certificate and a private key, together. The certificate identifies the user, and the private key corresponds to the certificate. A single user can have several identities; for any given user each certificate could have a different name, email address, or issuer. These identities would be used for different security contexts. For example, one could be used to sign others' certificates, and one could be used to identify the user by email, and these need not be the same identity.

In the context of the Mac OS X Server Certificate Manager, identities 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.

About 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.

About Intermediate Trust

If you are your own Certificate Authority, and your certificates are not trusted by the default shipping root certificates in Mac OS X, your clients can still be configured to trust your certificates through intermediate trust.

Trust is the ability of a client to believe the identity of a server when it connects. A trusted server is a known server that the client can transact with securely, without interference from outside and unknown parties.

Mac OS X clients follow x.509 trust validation when accepting certificates, meaning they follow the chain of certificate signers back until they find a trusted self-signed root certificate.

Often times, the longer the chain of trust, the more certificates out in the world have been signed, and the greater the chance that any given certificate can't actually be trusted. In order to deal with this scope issue, Mac OS X lets you specify a trusted anchor (in other words, a certificate which is not a root CA self-signed certificate, but you trust nonetheless). A client can trust a certificate closer in the chain of trust, or even just the submitted certificate itself. Trusting a certificate which isn't a shipping root anchor is intermediate trust.

To accomplish this, trust needs to be bestowed on certificates instead to keychains (as was done previously). In v10.4, trust was automatically given to certificates in the keychain called "X509Anchors." The X509Anchors keychain was deprecated starting with Leopard v10.5.

There are several keychains that can hold certificates:

- **SystemRootCertificates:** This keychain holds all the root certificates that ship with Mac OS X. The certificates there have already had trust given to them.
- **System:** This keychain holds all the certificates that the computer administrator can add. All the users on a given client can read from this keychain. The trust settings of a certificate in this keychain can override those of a certificate in SystemRootCertificates.
- **Any other keychain:** This holds all the certificates for a given user, and is only accessible to that user. The trust settings of a certificate in this keychain can override those of a certificate in SystemRootCertificates or System.

Trusted certificates can be in any of these locations, but in order to trust a certificate, trust settings must be given explicitly to the desired certificate.

To configure your clients to trust your 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 Drag the certificate to the System keychain using Keychain Access. Authenticate as an administrator, if requested.
- 4 Double click the certificate to get the certificate details.
- 5 In the details window, click the Trust disclosure triangle.
- 6 Select "Always Trust" from the pop-up menu next to "When using this certificate."

You have now added trust to this certificate, regardless of who it is signed by.

From the command-line

After copying the certificate over to the target client computer, perform the following where <certificate> is the file path to the certificate:

```
sudo /usr/bin/security add-trusted-cert -d -k /Library/Keychains/System.  
keychain <certificate>
```

You can use the security tool to save and restore trust settings as well. For more information on using the `security` command-line tool, see the `security` man page.

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.

Image
file:///Volumes/ServerAdmin/documents/Content/Art/
509_scert_010-0932.a1
not found.

Certificate Manager provides integrated management of SSL certificates in Mac OS X Server for all services that allow the use of SSL certificates. On installation, the server creates a self-signed certificate for immediate use from information you put in during server setup.

Certificate Manager uses Mac OS X's Certificate Assistant 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 then accessible by the 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 Certificate Assistant in Keychain Access (located in /Applications/Utilities/). It provides these capabilities and others for working with x.509 certificates.

Identities that were created and stored in OpenSSL files can also be imported into Certificate Manager. They are accessible to services that support SSL. 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 expiration date of the certificate
- When selected, the detailed contents of the certificate

When certificates and keys are imported via Certificate Manager, they are put in the `/etc/certificates/` directory. The directory contains four PEM formatted files for every identity:

- The certificate
- The public key
- The trust chain
- The concatenated version of the certificate plus the trust chain (for use with some services)

The certificate and trust chain are owned by root user and wheel group, permissions set to 644. The public key and concatenation file are owned by root user and certusers group, permissions set to 640.

Each file has the following naming convention:

`<common name>.<SHA1 hash of the certificate>.<cert | chain | concat | key>.pem`

For example, the certificate for a web server at `example.com` might look like this:

`www.example.com.C42504D03B3D70F551A3C982CFA315595831A2E3.cert.pem`

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, create a certificate then generate a Certificate Signing Request (CSR) to send to a CA, or import certificates previously created with OpenSSL.

If you have previously generated certificates for SSL, you can import them for use by Snow Leopard Server services. The OpenSSL keys and certificates need to be in PEM format.

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.

When you set up Mac OS X Server, the Server Assistant creates a self-signed certificate based on information you provided on setup when it's first installed. It can be used for any service that supports SSL. When your clients choose to trust the certificate, SSL connections can be used without user interaction from that point on.

This initial self-signed certificate is used by Server Admin and Server Preferences automatically to encrypt administrative functions.

Creating a Self-Signed Certificate

A self-signed certificate is generated automatically at server setup. Although it is available for use, you may want to customize the information contained in the certificate, so you would create a new self-signed certificate. This is especially important if you plan on having a Certificate Authority sign your certificate.

When you create 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.

If you're using a self-signed certificate, consider using intermediate trust for it and import the certificate into the System keychain on all client computers (if you have control of the computers). For more information about using intermediate trust, see “About Intermediate Trust.”

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 and choose Create a Certificate Identity.

Certificate Assistant launches, pre-populated with some needed information to generate the certificate.

- 4 If you choose to override the defaults, check “Let me override defaults” and follow the onscreen instructions.
- 5 When finished, click Continue.
- 6 Confirm the certificate creation by clicking Continue.

The Certificate Assistant generates a key pair and certificate. Certificate Manager encrypts the files with a random passphrase. It puts the passphrase in the System keychain, and puts the resulting PEM files in `/etc/certificates/`.

Requesting a Certificate From a Certificate Authority

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

You will need a certificate for the CA to sign. You can use the one which was generated at server setup, but more likely you will want to generate one that has all the details the CA requires before signing. If you need to generate a certificate before getting it signed, see “Creating a Self-Signed Certificate.”

To request a signed certificate:

- 1 In Server Admin, select the server that has services that support SSL.
- 2 Click Certificates.

- 3 Select the certificate you wish to have signed
- 4 Click the Action button below the certificates list, and choose “Generate Certificate Signing Request (CSR).”

Certificate manager creates the signing request, and shows the ASCII text version in the sheet.

- 5 Click Save to save the CSR to the disk.

Your Certificate Authority will have instructions on how to transfer the CSR to the signer. Some CAs require you to use a web interface, others require sending the CSR in the body of an email message. Follow the instructions given by the CA.

The CA will return a digital signature for you to add to your certificate. For instructions on what to do now with your completed signature, see “Replacing an Existing Certificate.”

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 or anchor 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 Identity

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.

The OpenSSL keys and certificates need to be in PEM format.

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 Add (+) button, and choose Import a Certificate Identity.
- 4 Drag the PEM file containing the private key to the sheet.
- 5 Drag the PEM file containing the public certificate to the sheet.
- 6 If desired, drag any associated non-identity certificates to the sheet as well.
- 7 Click the Import button.

If prompted, enter the private key passphrase.

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.

If you have previously generated certificates for SSL, you can import them for use by Snow Leopard Server services. The OpenSSL keys and certificates need to be in PEM format.

If you chose custom locations for your SSL certificates with Leopard Server, you will need to import them into Certificate Manager if you want them to be available for Snow Leopard Server services. Custom filesystem locations for certificates cannot be managed for services using Server Admin for Snow Leopard. If you want to continue to use custom file locations, you will need to edit the configuration files directly.

When certificates and keys are imported via Certificate Manager, they are put in the `/etc/certificates/` directory. The directory contains four PEM formatted files for every identity:

- The certificate
- The public key
- The trust chain
- The concatenated version of the certificate plus the trust chain (for use with some services)

Each file has the following naming convention:

`<common name>.<SHA1 hash of the certificate>.<cert | chain | concat | key>.pem`

For example, the certificate for a web server at example.com might look like this:

`www.example.com.C42504D03B3D70F551A3C982CFA315595831A2E3.cert.pem`

Once imported, Certificate Manager encrypts the files with a random passphrase. It puts the passphrase in the System keychain, and puts the resulting PEM files in `/etc/certificates/`.

Editing a Certificate

After you add a certificate signature, you can't edit the certificate. You must replace it with one generated from the same private key.

For instructions on how to do this, see “Replacing an Existing Certificate.”

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 your certificate to your clients:

- 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 Drag the certificate to the System keychain using Keychain Access. Authenticate as an administrator, if requested.
- 4 Double click the certificate to get the certificate details.
- 5 In the details window, click the Trust disclosure triangle.
- 6 Select “Always Trust” from the pop-up menu next to “When using this certificate.” You have now added trust to this certificate, regardless of who it is signed by.

From the command-line

After copying the certificate over to the target client computer, perform the following where `<certificate>` is the file path to the certificate:

```
sudo /usr/bin/security add-trusted-cert -d -k /Library/Keychains/System.  
keychain <certificate>
```

You can use the `security` tool to save and restore trust settings as well. For more information on using the `security` command-line tool, see the `security` man page.

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. Renewing a certificate is the same as replacing a certificate with a newly generated one with an updated expiration date.

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 renew.
- 5 Click the action button and select “Replace Certificate with Signed or Renewed Certificate.”
- 6 Drag the renewed certificate to the sheet.
- 7 Click Replace Certificate.

Replacing an Existing Certificate

If you change the DNS name of the server or any virtual hosts on the server, you will need to replace an existing certificate with an updated one.

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 replace.
- 5 Click the action button and select “Replace Certificate with Signed or Renewed Certificate.”

- 6 Drag the replacement certificate to the sheet.
- 7 Click Replace Certificate.

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. Notably, portable home directory synchronization is provided via SSH.

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) {
\Qssh $server -x -o batchmode=yes shutdown -r now\Q
}
}
```

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

Mac OS X Server can use another level of access control for added security. Administrators can be limited as to which service 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 their assigned services. This results in a kind of access control for individual server features and services.

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:
 - Select “Allow only users and groups below.”
 - Click the Add (+) button to open the Users & Groups window.
 - Drag users and groups from the Users & Groups window 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 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 78 .
- 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 Snow Leopard Server on a computer with Mac OS X makes the computer a server with Mac OS X Server.

Installing Snow Leopard Server on an existing server with an Mac OS X Server v10.2–10.5 upgrades the server software to v10.6 Snow Leopard Server.

This chapter includes instructions for a fresh installation of Snow 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:

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:

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:

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:

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:

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:

Step 7: Set Up Services Restart from the target disk to proceed to setup. For more information about server setup, see Chapter 6, “Initial Server Setup,” on page 108.

System Requirements for Installing Mac OS X Server

The Mac desktop computer or server where you install Mac OS X Server v10.6 Snow Leopard must have:

- An Intel processor, 1.5 GHz or faster
- 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 Server 10.6

The serial number used with any previous version of Mac OS X Server will not allow registration in v10.6.

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 126.

If you're using an installation disc for Mac OS X Server v10.6, you can control installation from another computer using VNC viewer software. Open source VNC viewer software is available. Apple Remote Desktop, described on “Apple Remote Desktop” (page 49), 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 *Quick Start* that came with your Xserve, where these procedures are documented.

Gathering the Information You Need

Use the *Installation and Setup Worksheet* to record information for each server you want to install. The information below provides supplemental explanations for items on the worksheet.

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.
- **Static IP Address:** It is strongly recommended that you have a static IP address already planned and assigned to the server.
- **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.

If you want to use the server connected to another directory server (Open Directory, Active Directory, or other OpenLDAP), make sure you have the directory server DNS name and IP address before installation.

SSH Service During Installation

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 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 iCal Service Utility, 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 contains installers for all the administration tools for use with Mac OS X Server v10.6. This disc also contains installers for a few developer tools: PackageMaker and Property List Editor.

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 v10.6 Snow Leopard or Mac OS X Snow 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 on a computer with Snow Leopard or Snow 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 v10.6 Snow 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 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 the server. 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 118.

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 Screen Sharing or 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 v10.6 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 find servers that are awaiting install finding the the Bonjour service name "_sa-rspndr._tcp.". You can use the `dns-sd` tool to identify computers on the local subnet where you can install server software. Enter the following from a computer on the same local network as the server:

```
dns-sd -B _sa-rspndr._tcp.
```

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.

For more information about this password, see "About Server Serial Numbers for Default Installation Passwords."

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

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.

To access the computer using Screen Sharing:

1 Locate and select the server in the Shared section of a Finder window sidebar.

If the remote server isn't listed in the Shared section of a Finder window sidebar, you can connect by choosing Go > Connect to Server and then entering `vnc://serveraddress` where *serveraddress* is the DNS name or IP address of the server whose screen you want to share.

2 Select the remote server, and click Share Screen in the Finder window.

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

Don't specify a user name.

For more information about this password, see "About Server Serial Numbers for Default Installation Passwords."

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.

About Server Serial Numbers for Default Installation Passwords

Server serial numbers are used for more than just inventory tracking. The server's built-in hardware serial number is used as the default password for remote installation.

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 have had a main logic board replaced on an Intel Xserve, the built-in hardware password is "System S" (no quotes).

If you don't know the IP address and the remote server is on the local subnet, you can find servers that are awaiting install finding the the Bonjour service name "_sa-rspnдр._tcp.". You can use the `dns-sd` tool to identify computers on the local subnet where you can install server software. Enter the following from a computer on the same local network as the server:

```
dns-sd -B _sa-rspnдр._tcp.
```

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 *User's Guide* or *Quick Start* 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. See "To create an image of the Install DVD."

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. See “To restore the image to a free volume.”

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 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.

From the command-line

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, the first command creates a disk image “Installer.dmg” from the device at `disk1s1`. The second command scans the image “Installer.dmg” and readies it for restore.

```
hdiutil create -srcdevice disk1s1 Installer.dmg
asr imagescan --source Installer.dmg
```

To restore the image to a free volume

- 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 the Restore tab.
- 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
```

From the command-line

If you prefer to use the command-line, you use the `asr` tool to restore the image to the partition. Restoring the disk image to the partition will erase all existing data on the partition. The basic syntax is:

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

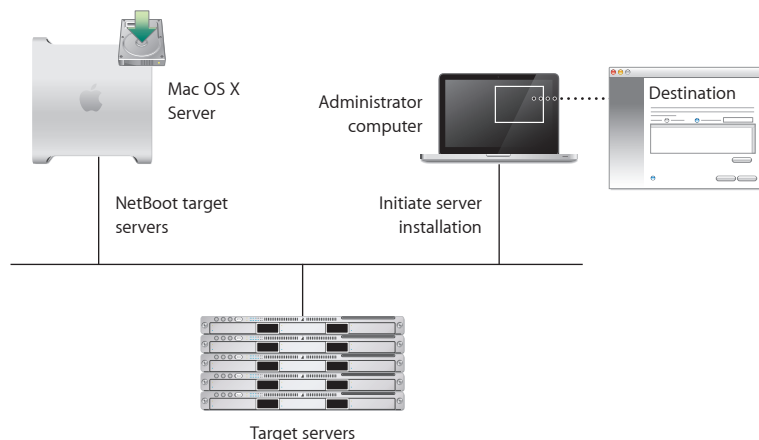
- ▶ **Tip:** You can use `asr` to restore a disk over a network, multicasting the blocks to client computers. Using the multicast server feature of `asr`, you could put a copy of the installer image on a partition of all computers that can receive the multicast packets. The `asr` tool can also fetch 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. For more information about `asr` and its capabilities, see the `asr` man page.

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

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

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.

Step 2: Start up the computer from the NetBoot server

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

Creating a NetInstall image from the 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*.

Starting up the computer from the NetBoot server

- 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 v10.6, 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
- Mac OS Extended (Journaled, Case-Sensitive) aka HFSX

About 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

About 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.

About Hard Disk Partitioning

The minimum recommended size for an installation partition is 20 GB. A much larger volume is recommended for a any configuration that keeps shared folders and group websites on the startup volume together with the server software.

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.

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.

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

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.

Partitioning a Disk 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 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 Mac computer with Mac OS X v10.6 and choose `Help > Disk Utility Help`.

From 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 Intel-based Mac computer.

The basic syntax is:

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

So the command is:

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

About 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 90.

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 Mac computer with Mac OS X v10.6 and choose Help > Disk Utility Help.

From 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 Mac computer with Mac OS X v10.6 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, DNS 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 v10.6 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 find servers that are awaiting install finding the the Bonjour service name "_sa-rspndr._tcp.". You can use the `dns-sd` tool to identify computers on the local subnet where you can install server software. Enter the following from a computer on the same local network as the server:

```
dns-sd -B _sa-rspndr._tcp.
```

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.

Similarly, servers awaiting setup use the service name "_svr-unconfig._tcp.", and can be found by entering:

```
dns-sd -B _svr-unconfig._tcp.
```

- **The preset password for the target server:** 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 have had a main logic board replaced on an Intel Xserve, the built-in hardware password is "System S" (no quotes).

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 Intel-based 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 90.

- 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 previous versions of Mac OS X Server 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 119.

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

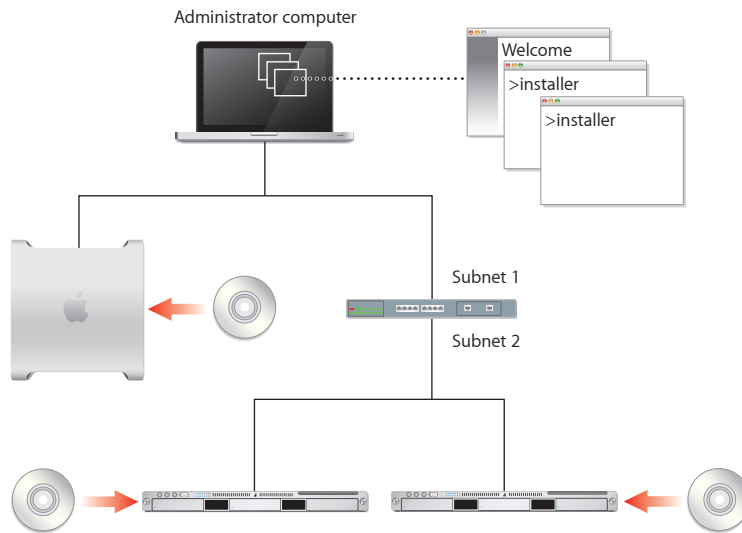
After installation is complete, the computer restarts and you can perform initial server setup.

- 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:
 - Quit Server Assistant when it starts automatically on the administrator computer.
 - Shut down the administrator computer and the server.
 - 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.

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 v10.5 Leopard or earlier, do not use them with Snow 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 Screen Sharing and 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 Admin in the /Applications/Server/ folder on the administrator computer.
- 2 Select the target server from the list of servers waiting for installation.

If neither the target server nor the list appear, make sure the target server is on the same local subnet as the administrator computer.

You can add the server manually by selecting Install Remote Server from the File Menu of Server Admin.

- 3 If the target server is 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).

- 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 your Intel Xserve has had its main logic board replaced, use "System S" (no quotes).

- 5 Proceed by following the onscreen instructions.

- 6 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 previous versions of Mac OS X Server 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 119

- 7 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 108 describes how.

Installing Remotely with Screen Sharing and VNC

If you're using an installation disc for Mac OS X Server v10.6 or later, you can control installation from another computer using a VNC viewer, like Mac OS X's built-in Screen Sharing, 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 detailed instructions for connecting to a computer running from an Install DVD, see “Remotely Accessing the Install DVD” on page 83.

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

Changing a Remote Computer’s Startup Disk

Sometimes you may need to explicitly set a remote computer’s startup disk. You can do this via the command-line using the `bless` command.

To change a remote computer’s startup disk

```
# Method 1
sudo bless --folder "/Volumes/<disk>/System/Library/CoreServices"
        --setBoot
sudo shutdown -r now
# Method 2
sudo systemsetup -liststartupdisks
sudo systemsetup -setstartupdisk <path to disk root>
```

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 the `installer` man page.

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 find servers that are awaiting install finding the the Bonjour service name "_sa-rspndr._tcp.". You can use the `dns-sd` tool to identify computers on the local subnet where you can install server software. Enter the following from a computer on the same local network as the server:

```
dns-sd -B _sa-rspndr._tcp.
```

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/ServerSetup/SerialNumber`.

If you're installing on an older computer that has no built-in hardware serial number, use 12345678 for the password. If you're installing on an Intel Xserve that has had its main logic board replaced, use "System S" (no quotes).

- 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/ServerNetworkImage10.6/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 90.

If the target volume has the latest Mac OS X Server v10.5 or 10.4.11 installed, when you run `installer` it will upgrade the server to v10.6 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 Chapter 6, "Initial Server Setup," on page 108.

Installing Multiple Servers

You can use VNC viewer software, or the `installer` tool to initiate multiple server software installations.

After using a VNC viewer to control installation of Mac OS X Server v10.6 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.

Additionally, you can create a NetInstall server on the target servers' local network which can automatically install an operating system. If you combine this with saved auto setup files, you can easily automate installation of multiple computers without much human interaction.

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 v10.6 to upgrade a desktop computer that has the following characteristics:

- Has Mac OS X v10.6 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 80

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 automatically 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.

Note: The Server Updates button refers only to updates for the server's operating system software from Apple. Third-party software is not updated when used. Additionally, it does not control software updates hosted in the Software Update service.

- Use the Software Update pane of System Preferences, if you are logged locally into the server.
- Use the `softwareupdate` command-line tool.
- Download a disk image of the software update from:
www.apple.com/support/downloads

Initial Server Setup

6

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.

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 *Installation and Setup Worksheet*. The following information provides supplemental explanations for some items on the worksheet.

When you're upgrading from the latest Mac OS X Server v10.5 or v10.4.11, Server Assistant displays the existing server settings, but you can change them. Use the *Installation and Setup Worksheet* to record settings you want the v10.6 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:

```
sudo shutdown now
```

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 DNS name, stored in directories and configuration files on the server, must be updated. For more information, see “Changing the Server's DNS 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 *Installation and Setup Worksheet*.

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.
- 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 (in `/System/Library/CoreServices/`) to manage Open Directory settings.

Specifying Initial Open Directory Usage

During setup of Snow Leopard Server, 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 (in `/System/Library/CoreServices/`), 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 v10.4.11 or the latest v10.5.
- **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 v10.6 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 v10.6. Open Directory in Mac OS X Server v10.6 supports MSCHAPv2 authentication, but Password Server in Mac OS X Server v10.2 doesn't support MS-CHAPv2.

After setup, use the Directory Utility (in `/System/Library/CoreServices/`) 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 v10.6 from the latest v10.5 or 10.4.11, 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 (in `/System/Library/CoreServices/`) 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 (in `/System/Library/CoreServices/`).

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 v10.4 or earlier, you may not be able to take advantage of some features:

- iCal Server requires an Open Directory server of v10.5 or later.

- In v10.3 and earlier, trusted directory binding, support for LDAP subdomains, and Directory Utility controls aren't available.
- Kerberos configuration is much more complex in v10.2. In addition, automatic synchronization of Kerberos and Password Server requires v10.3 or later.

Binding a Server to Multiple Directory Servers

Automatic server setup allows you to bind to multiple servers. However, the option is not present in the Server Assistant, so you will not be able to do it interactively. You need to save setup data, then you have to modify the plist by hand

In the saved setup data, you will find the "directoryServers" key in the plist, and it's an array. You add items (or in this case directory servers) to the array, the server binds to all of them.

For more information on making saved server setup data, see "Using Automatic Server Setup" and "Creating and Saving Setup Data."


Using Interactive Server Setup

The simplest way to set up a small number of servers is to use Server Admin'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 Admin and choose Server > Configure Remote Server, connect to one or more target servers, supply setup data, and then initiate the setup immediately.

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 Intel 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 116.



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 115.

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.

To set up a local server interactively:

- 1 Fill out the *Installation and Setup Worksheet*.

Supplemental information appears in “Information You Need” on page 108.

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), click Save As.
To encrypt a configuration 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.
- 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 *Installation and Setup Worksheet*.
Supplemental information appears in “Information You Need” on page 108.
- 2 Make sure the target server is running.
- 3 On an administrator computer, open Server Admin.
- 4 In the Server menu, select “Configure Remote Server.”
The Server Assistant launches.
- 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 In the Language pane, specify the language you want to use to administer the target server.
- 7 If you are using saved setup data, do the following:
 - In the Language pane, choose File > Open Configuration File 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.

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 DHCP or DNS servers you specify for the server you're setting up are running.
- 8 After you specify setup data, review the summary displayed by Server Assistant and optionally click Go Back to change data.
 - 9 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), click Save As.
To encrypt a configuration 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.
 - 10 To initiate setup of the remote target server, click Apply.
 - 11 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, they must all be Intel processors)
- The servers use the same setup data, except for server software serial numbers and network identities (DNS 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 *Installation and Setup Worksheet* with the settings you want to use for all servers you want to set up.

Supplemental information appears in "Information You Need" on page 108.

- 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 Admin.
- 4 In the Server menu, select “Configure Remote Server.”
The Server Assistant launches.
- 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 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), click Save As.
To encrypt a configuration file, 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 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 or removable storage. 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.

- “Creating and Saving Setup Data” on page 120
- “Using Encryption with Setup Data Files”

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 119
- “Setting Up Servers Automatically Using Data Saved in a File” on page 121

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.

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.

- First, looks on all volumes for a folder at the root named "Auto Server Setup," starting at the boot volume and then the rest alphabetically.
- It searches through "Auto Server Setup" folders, looking for a file with the extension ".plist". There is no naming convention for the plist.
The plist file must contain the key "VersionNumber" with value "4" or it will be ignored.
- It evaluates all the profile plists found to evaluate the most specific match.

Most specific to least specific criteria are:

- hardware serial number
- MAC address
- IP address
- DNS name (fully qualified)
- Computer name
- none of the above

If a saved setup data profile contains multiple network connection services, Server Assistant first tries to match hardware (MAC) addresses, and failing that try to match interface (BSD port) names. If a profile has multiple conditions, it applies to a computer that satisfies any of them

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.

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.

The best way to create setup data is to use Server Admin to launch the Server Assistant, which lets you work with setup data without connecting to specific servers. You specify setup data, then save it in a file. Target servers where Mac OS X Server v10.6 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.

To create a setup data file:

- 1 For each server to set up, fill out the *Installation and Setup Worksheet*.
- 2 On an administrator computer, open Server Admin.
- 3 In the Server menu, select "Create Auto Server Setup File."
The Server Assistant launches.
- 4 In the Welcome pane, select "Save 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 servers.
- 6 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.
- 7 Click Continue and enter the setup data as you move through the Assistant panes, following the onscreen instructions.
- 8 After you specify setup data, review the summary displayed by Server Assistant and optionally click Go Back to change data.

Use the Save Configuration Pane after entering setup information to save it as an auto setup profile. When you click Save, you can give the profile any filename you like as long as it ends with .plist. You can save a variation of the profile by clicking Go Back to change some settings, clicking Continue to return to this pane, and clicking Save again.

- 9 Click Save As, then select Configuration File.

- 10 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.

Using Encryption with Setup Data Files

Saved setup data can be 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 with the setup data of the target server. If you want to store the password for non-interactive setup, the passphrase should be in a plain text file with the extension “.pass” in the same folder as the setup data for the server.

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.

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 installation.

To apply setup data from a file:

- 1 For each server to set up, fill out the *Installation and Setup Worksheet*.
- 2 If you want to have the server configure itself without further input, place the previously generated auto setup data file in a location where target servers can detect it.

For more information on where the server looks for setup data, see “How a Server Searches for Saved Setup Data.”

If you have not previously created saved setup data, see “Creating and Saving Setup Data.”

If the setup data is encrypted, make the passphrase available to target servers. For more information, see “Using Encryption with Setup Data Files.”

- 3 On an administrator computer, open Server Admin.
- 4 In the Server menu, select “Configure Remote Server.”
The Server Assistant launches.
- 5 In the Language pane, select “Open Configuration File.”
- 6 Navigate to the location where you saved the saved setup data file.
- 7 If the setup data is encrypted, supply the passphrase.

- 8 To provide a passphrase interactively to a remote server which already has its setup data, use Server Assistant on an administrator computer that can connect with the target server.
 - In the Welcome or Destination pane, choose File > Supply Passphrase.
 - In the dialog box, enter the target server's IP address, password, and passphrase.
 - Click Send.
- 9 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`.

Setting a Mac OS X Server Serial Number with the Command-line

After an automatic setup, you may need to set a specific Mac OS X serial number for your server. For example, you may have completed an automatic setup with a generic setup data profile, and now need to put individual serial numbers to their respective servers.

To set the server serial number

```
sudo serversetup -setServerSerialNumber <serialnumber-dash-separated>
    [<name> <organization>]
```

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, 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 in the same folder as the setup data.

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 installation and initial startup, you 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 All Other Services

All the services of Mac OS X Server Snow Leopard require specialized setup instructions in order to tailor the service to your specific needs.

See the individual service guides for step-by-step instructions for setting up and managing the services.

This chapter shows you how to complete ongoing management for your systems, including setting up administrator computers, designating administrators, and maintaining service uptime.

You should read the following sections as a basic introduction to Mac OS X Server management:

- “Computers You Can Use to Administer a Server” on page 126
- “Using the Administration Tools” on page 127
- “Changing the Server’s Computer Name and the Local Hostname” on page 146
- “Adding and Removing Servers in Server Admin” on page 130
- “Administering Services” on page 147
- “Tiered Administration Permissions” on page 151
- “Workgroup Manager Basics” on page 152

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 v10.6 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 126.

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 v10.6 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 v10.6 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 v10.4 or 10.5.	"Setting Up Users' Macs Automatically" (page ###)
Server Assistant	Set up a v10.6 server.	"Using Interactive Server Setup" (page 113)
Workgroup Manager	Administer accounts and their managed preferences.	"Workgroup Manager Basics" (page 152)
Server Admin	Configure and monitor services and administrator access, and configure share points. Set up and manage QuickTime media streaming.	"Working with Settings for a Specific Server" (page 132) "Server Admin" (page 38)
Gateway Setup Assistant	Configure a simple gateway between two networks	
iCal Service Utility	Add locations and resources to your iCal server.	"iCal Service Utility"
System image tools	Manage NetBoot and NetInstall disk images.	"System Image Management" (page 47)
Server Monitor	Monitor Xserve hardware.	"Server Monitor" (page 174)
QTSS Web Admin, QuickTime Broadcaster, and QuickTime Player	Manage media playlists and prepare it for streaming or progressive download.	"Media Streaming Management" (page 47)
Apple Remote Desktop (optional)	Monitor and control other Macintosh computers.	"Apple Remote Desktop" (page 49)
Command-line tools	Administer a server using a UNIX command shell.	"Command-Line Tools" (page 48)
Xgrid Admin	Monitor local or remote Xgrid controllers, grids, and jobs.	"Xgrid Admin" (page 49)

Working With Pre-Version 10.6 Computers From Version 10.6 Servers

You can use the version of Server Admin included with Mac OS X Server v10.6 to administer the latest Mac OS X Server v10.5.

Using Mac OS X Server v10.6 will not administer DNS hosted on a server version earlier than v10.6.

Workgroup Manager on a v10.6 server can be used to manage Mac OS X clients running the latest Mac OS X v10.5. However, after you edit a user record using Workgroup Manager on v10.6, you can only access it using Workgroup Manager on v10.6.

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:

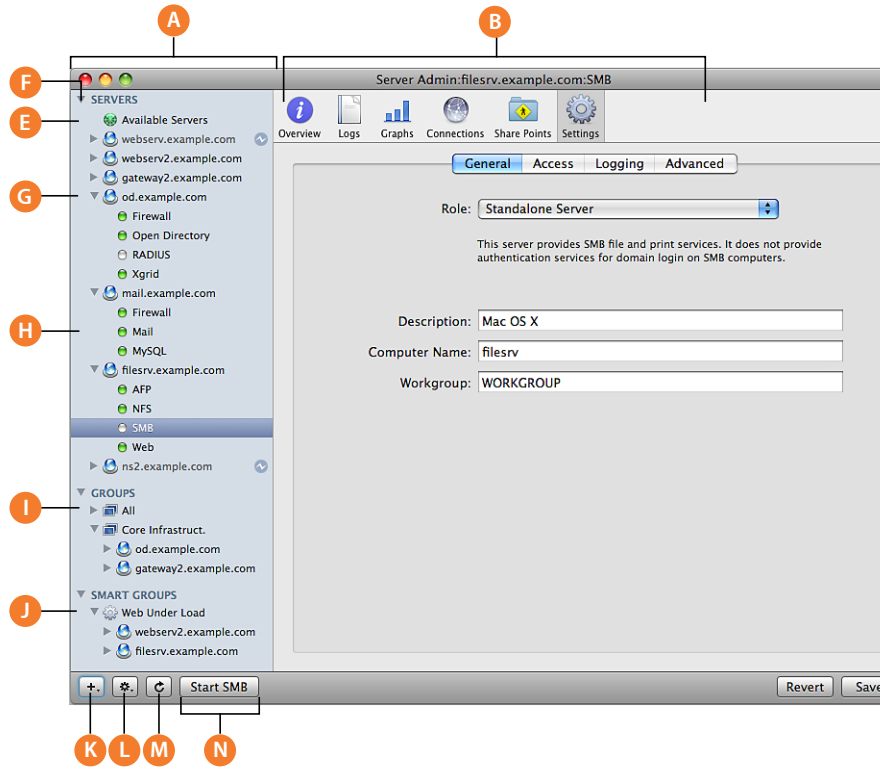
Port number and type	Service
22 TCP	SSH command-line shell
311 TCP	Server Admin (with SSL)
626 UDP	Serial number support
625 TCP	Remote Directory Access
ICMP incoming and outgoing	standard ping
53 UDP	DNS name resolution

Server Admin Basics

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.

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:

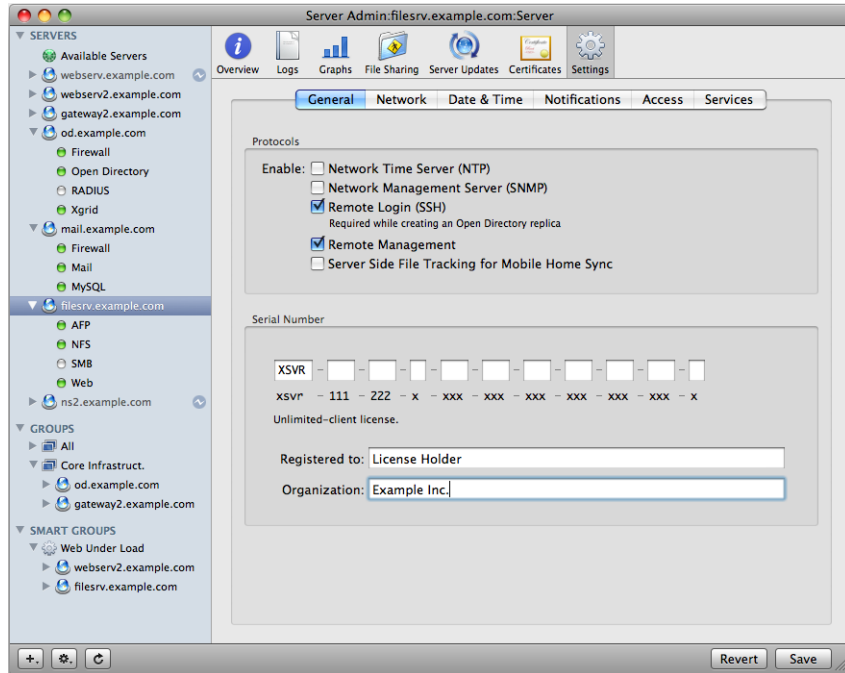
- 1 Click the Add (+) button under the Server list at the bottom of the Server Admin window.
- 2 Select Add Smart Group.
- 3 Name the smart group.
- 4 Define the criteria by which servers will appear in the list, and click OK.

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. This only controls updates to the server's own 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 "Server Status 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 152.
- **Services pane:** Click Services to show or hide services in Server Admin for this server.

Understanding Changes to the Server IP Address or Network Identity

When you change a server's IP address, DNS name, local hostname, or Computer name, there may be additional configuration steps needed for each service provided. Each service relies on IP addresses or names differently, therefore the exact combination of steps relies on your individual setup.

The following sections give guidance regarding what types of changes will be necessary in the event of a name or IP address change.

Understanding Mac OS X Server Names

There are three names used by Mac OS X Server: computer name, local hostname, and DNS name. They are used by different parts of the system for different reasons, and are not linked. Changing the computer name, and the local hostname is not the same thing as changing the DNS name.

The computer name is a user-friendly name for the system, and is shown in the Finder and tools like Apple Remote Desktop.

The local hostname is a domain name, usable only on the local network, and is published to other services which are Bonjour-aware.

The DNS name is the Internet host name, which is a fully qualified domain name. Only the DNS name is the Internet-routable name that services use for network identity.

Understanding IP Address or Network Identity Changes on Infrastructure Services

Certain services are infrastructure services. This means that they provide the basic addressing, name resolution, and routing necessary for other services to function. Infrastructure services include:

- DNS
- DHCP
- Directory Service
- Firewall
- Mobile Access
- NAT
- NetBoot
- RADIUS
- VPN

Generally, changing the IP address or name of an infrastructure server requires an intimate knowledge of the new network configuration and topology as well as lots of manual setting changes. Any changes to these infrastructure services can cause widespread disruption of other services until the correct setting modifications are made.

DNS

For a server not hosting DNS, changing a server's IP address requires changes to the data in the DNS server. Minimally, the server's NS, A, and PTR records need to be changed. Since the DNS information for the server is hosted elsewhere, those records will need to be updated manually on the DNS server. Your network configuration may have other domains, computers, and record types that are impacted by a server's IP address change (SRV records, for instance). These other records should be examined thoroughly after any change to a server's IP address.

If the server is itself a DNS server, the tool `changeip` can be used to automatically change the NS, A, and PTR records. Changing a DNS server's IP address directly impacts any client computer which utilizes the DNS server. For example, the DNS server's IP address could be provided to DHCP clients automatically, so all DHCP clients rely on the DNS server's correct IP address. All the DNS names for all the domains hosted by the DNS server have to be examined.

Because of DNS caching, many of the clients may not respond to the changes in the DNS system as quickly as needed. In order to expedite DNS server setting propagation, update all wireless access points, DHCP servers, manually configured IP address clients, and DHCP address clients by restarting them or renewing their DHCP leases.

In summary, all clients that refer to the DNS server's IP address for name resolution need to be updated to use the new IP address.

Changing a server's DNS name or domain impacts all the other services which rely on the server's domain name resolving correctly in DNS. The affected services include:

- Directory service
- Kerberos service and Kerberos Realm names
- WINS server names
- DHCP supplied search domains

DHCP

Changing the IP address of a DHCP server could possibly invalidate all of the the subnets and static IP addresses handled by the server. Additionally, it is possible that the change in IP address may result in unreachable search domain names, WINS server names, or LDAP URLs. Examine these settings, if needed.

Many of the clients may not respond to the changes in the DHCP system immediately. After a DHCP server change, update all wireless access points, manually configured IP address clients, and DHCP address clients by restarting them or renewing their DHCP leases.

Changing the domain name of a DHCP server could also obsolete the search domain names, WINS server names, or LDAP URLs. Changing the only hostname segment to a fully qualified domain name may not have the same effect.

Directory Service and Kerberos

Changing the IP address of an Open Directory Server may possibly invalidate some of the data within the records themselves (computer records, or user home directories). None of the contents of the records are altered automatically when you change the IP address, only the configuration.

Changing the DNS name of the directory server will require that all the bound machines will need to be rebound to the new directory name and address.

If you have set up a Kerberos environment, note that the Kerberos realm does not change when the hostname is changed.

Firewall

Changing the IP address of the Firewall can significantly alter the effectiveness of the service. In Mac OS X Server v10.6, IP firewall rules are stored and referenced as “address groups.” A change to the IP address of the firewall server may render traffic to the address groups unable to be routed, and therefore none of the specific firewall rules would be applied.

Check all the firewall rules when changing the IP address of the firewall server.

Mobile Access (Proxy Services)

Most proxy services should remain relatively unaffected by a change to IP address or domain name. If you have edited the `com.apple.securityproxy_mail.plist` manually to have the proxy server connect to *itself* for some service by some other address than the link-local address (127.0.0.1 or localhost), you will have to change it manually again.

However, proxy services are affected if the IP address or DNS name of the destination servers changes. Any time you change one of the *proxied* services’ names or addresses, you will need to reconfigure Proxy Service.

If you have configured an HTTP Secure Proxy virtual host, will need to delete and re-create the proxy mappings of any proxied servers.

NAT

Network Address Translation should not be affected by a change to the server’s IP address or DNS name. All of the clients behind the address translation still have contact with the NAT router by the internal IP address. If you have made any manual modifications to the NAT service configuration files, make sure those changes are compatible with the new IP address or DNS name.

Netboot

Netboot will not require any reconfiguration after changing the IP address or DNS name. However, all the clients which use it will need to reselect the server after the changes.

RADIUS

If you change the RADIUS server IP address, you might have to check or reconfigure the IP addresses of the associated base stations. Additionally, if you're using SSL certificates, you need to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure the service's new certificate.

VPN

VPN servers server IP address ranges to VPN clients; they also server DNS server names and search domains. Any of these can be affected by a change to the VPN server's IP address or domain name. Additionally, the VPN server contains routing definitions which are based on IP addresses. A change to the IP address may make those routing addresses unreachable.

Check all the VPN settings when changing the IP address of the VPN server.

Understanding IP Address or Network Identity Changes on Web and Wiki Services

Certain services can be classified as web services. This means that they provide the interaction, back end database storage, and media streaming of websites hosted on the server. Web services include:

- Web
- MySQL
- QTSS
- Wiki
- Certificates for web and wiki service

Generally, these services in the initial default configuration are resilient and automatically adjust to any changes made to either the IP address or the server name. However, your web services might be highly customized, and therefore need manual configuration changes to maintain service integrity.

Web

If you change the web server's DNS name or IP address, you will need to modify the domain name and web server aliases. You should also check the site load balancer members.

If you change the web server's DNS name, you will need to modify any virtual hosts that use SSL. All virtual hosts which use SSL will need new certificates. You may have to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure each virtual host's new name and certificate.

If you change the web server's IP address(es), you need to use Server Admin to change any virtual hosts that use a specific IP address. The default wild-card virtual host doesn't need to be modified.

For either change, if you configured Mobile Access for web (or possibly other proxy settings), you need to delete and recreate the proxy settings for all affected hosts.

MySQL

In general, MySQL would not be affected by changing an IP address or DNS name. However, none of the data within the databases is altered automatically when the DNS name or IP address are changed. You are responsible for replacing references to the DNS name and address (if even used) within your databases.

If you have set a database "root" password, there may be entries in the database GRANT table (database=mysql, table=user) that refer to the previous server DNS name. In this case, you should use Server Admin to reset the root password, which will then reflect the current server identity.

Server administrators should make sure that any MySQL clients that have saved references to the DNS name of the MySQL Service are updated to reflect any change in the server identity.

QTSS

The typical default configuration will not have any further configuration after changing the DNS name or IP address of a QuickTime Streaming Server. If you have configured specific IP bindings, you will need to change those to the new address, and restart the service. Any relays you have defined may have invalid IP addresses after an IP address change.

Wiki

Wiki service remains unaffected by a change in the IP address, assuming apache is still functioning, and DNS name changes.

However, wikis can be configured to restrict to certain DNS names; if you have manually edited the configuration files in such a way as to restrict wiki access to certain DNS names, you will have to make the appropriate changes in those file.

Certificates for web and wiki services

Web and wiki servers which use SSL will need new certificates. You may have to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure each service's, or site's, new certificate.

Understanding IP Address or Network Identity Changes on File Services

File service provide file storage and retrieval for network clients. File services include:

- AFP

- SMB
- NFS
- FTP

For the most part, changing the network address or DNS name of a file server has no internal affect on file services. The file service processes monitor the network interfaces for changes and adapt as necessary without administrator intervention. No further configuration is required.

There are a few places where you might need to change configuration settings:

- **SMB:** The computer name defaults to the the unqualified primary DNS name. Changing the DNS name of the server will cause a mismatch between the DNS name and the defined “computer name.”
- **FTP:** The service can use SSL certificates, and will need new certificates. You may have to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure each service’s new certificate.

Additionally, clients might have URLs, bookmarks, “favorites” or documentation which refer to previous DNS names or IP addresses. You should make an effort to ensure that client information is updated to reflect the new name or address.

Finally, you might have other software which interface with file servers (for example, automated scripts) and refer to old DNS names and IP addresses. You should update those applications or scripts as well.

Understanding IP Address or Network Identity Changes on Mail Services

Mail services are the suite of services which provide email delivery, retrieval, and processing. Email services include:

- SMTP
- IMAP and POP
- Mailing List
- Anti-virus and anti-spam
- Certificates for mail services

Most of the mail services require some sort of restart after changing a DNS name or IP address of the mail server. If you have manually changed any of the configuration files, you may need to edit them manually again. Additionally, some mail services require a full shutdown and startup (rather than a simple service reload) to get the address and identity changes.

There are many places in the mail services configuration panes where you enter domain names, mail host names, relay host names, and email addresses. Any change you make to the DNS name could potentially have an affect on the service. Double check all name and IP address settings carefully.

IMAP and POP

Dovecot, the IMAP and POP service, loads the fully-qualified domain name at startup and configuration reload. After a change, Dovecot needs to be restarted or given a SIGHUP command, at a minimum).

You will also need to restart if you have manually edited any of the following configuration parameters: `listen`, or `ssl_listen`.

SMTP

Postfix, the SMTP service, is very sensitive to network address and identity changes. The information that it stores about the DNS name and the IP address and network interfaces is only loaded once at service startup. In order to resume service after a change to either the DNS name or the IP address, you need to fully stop the service, and restart it.

You will also need to restart if you have manually edited any of the following configuration parameters: `inet_interfaces`, `inet_protocols`, `smtp_bind_address`, `myhostname`, or `mydomain`.

Mailing List

Mailman, the mailing list service, keeps track of the incoming and outgoing mail hosts by reading them on startup. If you change the hostname or IP address you will need to restart Mailman for it to honor the configuration changes.

Anti-virus and Anti-spam

ClamAV, the anti-virus service, gets its listening address at startup as well. After making any changes to the DNS name or IP address, you will need to stop and restart to resume service.

SpamAssassin, the anti-spam service, gets its configuration information at startup and can have reload configuration data while running. To load new configuration data, SpamAssassin needs to be restarted or given a SIGHUP command, at a minimum.

Certificates for mail services

Mail servers which use SSL will need new certificates. You may have to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure each service's new certificate.

Understanding IP Address or Network Identity Changes on Collaboration Services

Collaboration services provide tools to coordinate people and time. Collaboration services include:

- Calendar Service
- Address Book Service
- iChat Service
- Certificates for collaboration services

Address Book Service

Changing the IP address of an Address Book server will not affect new connections to the server; however, it may disconnect existing client connections. If you have manually edited the `BindHTTPPorts` or `BindSSLPorts` options in the `carddav.plist` file, you will need to edit them again and restart the service.

Changing the DNS name of an Address Book server necessitates restarting the service. If you have manually edited the `ServerHostName` setting in the `carddav.plist` file, you may have to do so again before restarting the service.

iCal Service

The iCal Server is based on the same underlying technology as the Address Book Server, so the needs are the same.

Changing the IP address of an iCal server will not affect new connections to the server; however, it may disconnect existing client connections. If you have manually edited the `BindHTTPPorts` or `BindSSLPorts` options in the `caldav.plist` file, you will need to edit them again and restart the service.

Changing the DNS name of an iCal server necessitates restarting the service. If you have manually edited the `ServerHostName` setting in the `caldav.plist` file, you may have to do so again before restarting the service.

iChat Service

The iChat service itself is highly resilient to network and identity changes on the primary ethernet port. No additional configuration is necessary once you've changed the DNS name or IP address of the iChat server.

However, the jabber IDs associated with the server will not automatically update to the new iChat server DNS name. For example, changing the server from `example.com` to `example.net`, Joe's jabber ID (`joe@example.com`) doesn't automatically migrate to "`joe@example.net`."

The jabber IDs for all service users can be changed using the `jabber_autobuddy` tool. The tool modifies the database by changing the "@host.com" part of all users associated with the old domain to reflect the new domain, as well as secondary references (individual- and group-based buddies) that reference the old domain.

The following commands need to be run to migrate the jabber IDs:

```
sudo serveradmin stop jabber
sudo jabber_autobuddy --move-domain <old_domain> <new_domain>
sudo serveradmin start jabber
```

Additionally, the tool makes an automatic backup of the previous database (`/var/jabberd/sqlite/jabberd2_bak.db`), which can be stored or restored as needed.

Certificates for collaboration services

AddressBook, iCal, and iChat servers which use SSL will need new certificates. You may have to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure each service's new certificate.

Understanding IP Address or Network Identity Changes on Podcast Producer

Podcast Producer is a complex service; it uses a number of other services and computers to perform its work. Because several Xgrid Agent computers and camera capture computers depend on contact with the Podcast Producer server, any changes to IP address or DNS name need to be coordinated for all the affected computers, not just the main Podcast Capture server.

If Podcast Producer server is run on a computer providing DNS to a network, or is run on a computer providing directory services to a network as an Open Directory Master, resolve the conflicts and network conditions for those services before attempting to account for the changes done to Podcast Producer. For more information on how address and identity changes affect DNS and directory services, see "Understanding IP Address or Network Identity Changes on Infrastructure Services."

Changing the IP address or DNS name may necessitate changing settings for the following services and software:

- DNS server
- Open Directory server
- Xgrid Controller
- NFS file service
- XSAN and its MDC configuration (if used for file storage)
- Mail Services (if used by a workflow)
- Wiki Server (if used by a workflow)
- iChat Server (if used by a workflow)

- QMaster (if used by a workflow)
- Final Cut Server (if used by a workflow)

You can reduce the number of services you have to reconfigure by initially defining an alias record in the DNS (a CNAME record) and using the DNS name alias as the DNS name for configuration purposes.

If any of the listed servers use SSL, it will need new certificates. You may have to regenerate or repurchase the certificates. You will need to use Server Admin to import the new certificates, then configure each service's new certificate.

To change the IP address of the Podcast Producer computer:

- 1 Stop the Xgrid job queue when empty (or stop and empty it).
- 2 Reconfigure the DNS, Open Directory, DHCP, and other infrastructure services.
For example, in the DNS, change the "A" record IP address of the Podcast Producer server.
- 3 Use `changeip` to change the IP address of the Podcast Producer server.
- 4 Restart (or renew the DHCP leases of) all the Podcast Camera Agents.
- 5 Restart (or renew the DHCP leases of) all the Xgrid Agents used for the Podcast Producer workflow grid.

Alternatively, instead of restarting the computers, you can flush the Directory services cache (using `dscacheutil` and sending a HUP to the `mDNSResponder` daemon).

To change the DNS name of the Podcast Producer computer:

- 1 Stop the Xgrid job queue when empty (or stop and empty it).
- 2 Reconfigure the DNS, Open Directory, DHCP, and other infrastructure services.
For example, in the DNS, change the "A" record host name of the Podcast Producer server.
- 3 Use `changeip` to change the DNS name of the Podcast Producer server.
- 4 Restart (or renew the DHCP leases of) all the Podcast Camera Agents.
- 5 Restart (or renew the DHCP leases of) all the Xgrid Agents used for the Podcast Producer workflow grid.
Alternatively, instead of restarting the computers, you can flush the Directory services cache (using `dscacheutil` and sending a HUP to the `mDNSResponder` daemon).
- 6 Unbind the Podcast Camera Agents from the previous DNS name and rebind them to the new name.
- 7 Reconfigure all Xgrid Agents to use the new DNS name.
- 8 Reconfigure any services used in the workflow to reference the new DNS name, if needed.

- 9 Update and reissue an SSL certificates that contain the server's DNS name.
- 10 Reconfigure Kerberos service on the server (using Directory Binding or `ssoutil`).
- 11 Update any scripted or automated software that submits data to or polls data from Podcast Producer.

Understanding IP Address or Network Identity Changes on Other Services

The remaining services affected by changes to the IP address or network identity include:

- Print
- Software Update Service
- Xgrid

After a print of software update server changes DNS name or IP address, a number of the changes that need to be made by the clients. However, the following guidelines for the server should be followed.

Print

Print service needs no changes if the IP address changes. If the DNS name changes, the administrator needs to restart print service to re-register the service with Bonjour to publish the name change.

If you have made any custom configurations of the `cupsd.conf` file, or have configured `/Config/Printers` entries in the directory service, you should review those custom configurations and update them if needed.

If you have assigned per-queue printing quotas to user accounts, you should update the account quotas to reflect the new server DNS name if needed.

Also, make sure that any printing clients which have saved references to the DNS name of print queues are updated to use the new DNS name.

Software Update Server

Software Update service needs to be restarted after changes are made to either the DNS name or the IP address of the service. Afterward, you need to update the list of available software updates.

Also, make sure all clients that have saved references to the DNS name of the Software Update server are updated to use the new DNS name.

Xgrid

Xgrid service needs to be restarted after changes are made to either the DNS name or the IP address of the service. Any changes to the DNS name or IP address should be made when the Xgrid job queue is empty and stopped.

If you use Kerberos for client authentication to the controller, you need to resolve any Kerberos configuration before attempting to reconfigure for Xgrid service.

If you change the DNS name of the controller, reconfigure all the Xgrid Agents to use the new controller's new DNS name.

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.

Do not turn off the primary network interface, and then turn it back on with a different address than it had before. Several services will not get the needed notification in order to update their configuration.

Changing your IP address can have significant unintended consequences, depending on the services your server provides. For information on the effects of changing the IP addresses, see “Understanding Changes to the Server IP Address or Network Identity” on page 134.

The command-line tool `changeip` is a tool to accomplish manually what would normally be done automatically, and it is still available in Snow Leopard.

Changing the Server's DNS Name After Setup

If you change a server's DNS name after setup, the name must be changed with your DNS service provider.

Until the server's DNS name matches the name with the DNS service provider, several services will not function correctly. Changing your DNS name can have significant unintended consequences, depending on the services your server provides. For information on the effects of changing the DNS name, see “Understanding Changes to the Server IP Address or Network Identity” on page 134.

The DNS name is the Internet host name, which is a fully qualified domain name. Only the DNS name is the Internet-routable name that services use for network identity.

To change the DNS name

```
sudo scutil --set OldName <NewName.domain.tld>
```

The tool `scutil` can be used to set the computer name and the local hostname. For more information, see the `scutil` man page.

The command-line tool `changeip` is not supposed to be used to change DNS names, although the tool is still available in Snow Leopard.

Changing the Server's Computer Name and the Local Hostname

The computer name is a user-friendly name for the system, and is shown in the Finder and tools like Apple Remote Desktop.

The local hostname is a domain name, usable only on the local network, and is published to other services which are Bonjour-aware.

The tool `scutil` can be used to set the computer name and the local hostname. For more information about `scutil`, see the tool's man page.

The command-line tool `changeip` is not supposed to be used to change computer names.

To change the computer name and local hostname:

- Change the names in the Network pane of the Settings section for the server in Server Admin.

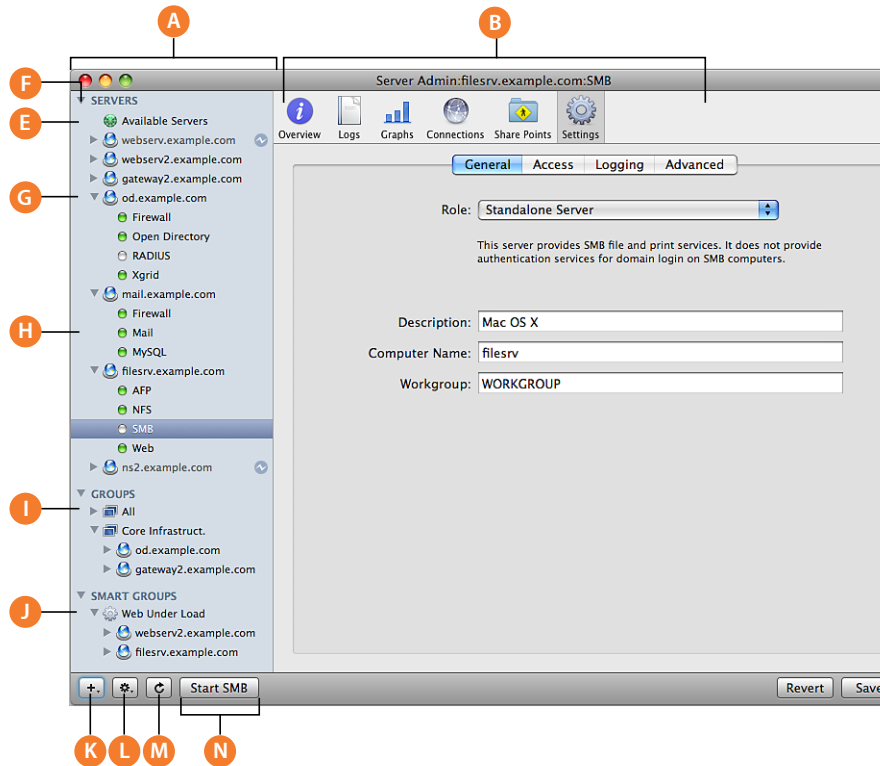
From the command line:

```
sudo scutil --set ComputerName <newComputerName>
sudo scutil --set LocalHostName <newLocalHostName>
```

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 service 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 service 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, 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. 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.

The following shows the Service Access Control List pane in Server Admin:

Image
"file:///Volumes/SrvrAdmin/documents/Content/Art/
500_saclaccess_019-0932.ai"
not found

To configure service access SACLs

- 1 Select a server in the Servers list.
- 2 Click Settings, then click Access.

- 3 Click Services.
- 4 Choose a service, and choose whether to allow everyone access to it, or whether to allow specified users access to the service.
- 5 If you have chosen to specify users, add the users and groups as desired.

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.

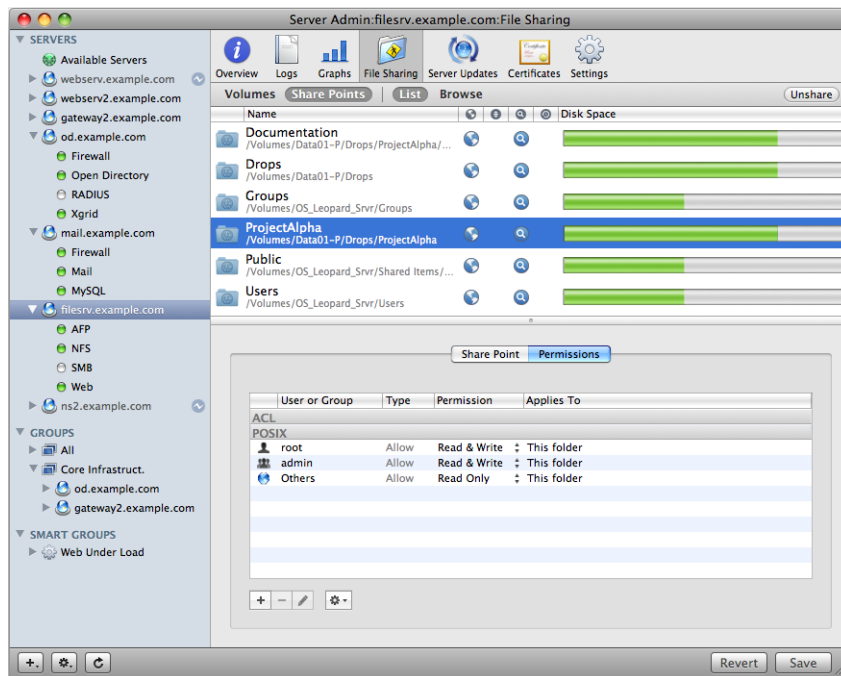
“Certificate Manager in Server Admin” on page 63 provides instructions for using Server Admin to create, organize, and use security certificates for SSL-enabled services. Individual service administration guides describe how to configure specific services to use SSL.

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 v10.6, 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. It is case-sensitive.

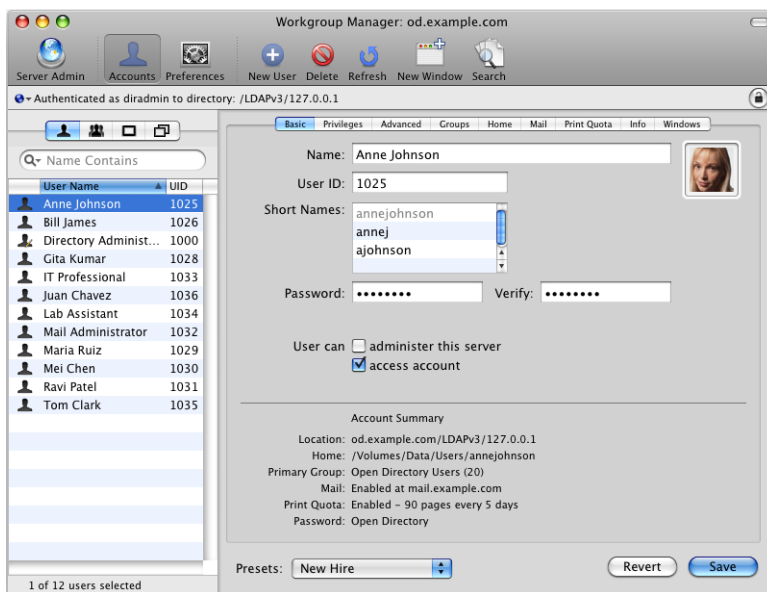
Administering Accounts

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 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:



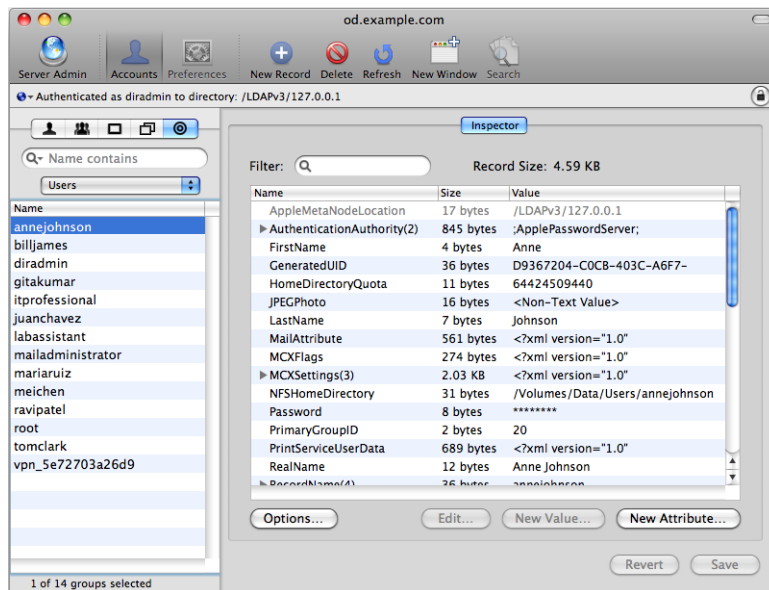
Click Details to use the preference editor to work with preference manifests. The following is a sample of the preference editor sheet in Workgroup Manager:

Image
 "file:///Volumes/SrvAdmin/documents/Content/Art/518_wgmprefeditor_019-0932.ai" not found.

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.

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:

- **Server Assistant:** This assistant helps you configure remote servers, install Mac OS X Server remotely, and make automatic server setup data files.
- **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 31.

The following is a list of configuration and data files for services available on Mac OS X Server.

Time Machine backs up service states and configuration files, but not files with your created data. To see which services Time Machine backs up, see “Understanding Time Machine as a Backup Tool” on page 36.

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

Address Book Service

File type	Location
Configuration files	/etc/cardavd/cardavd.plist
Data	/Library/AddressBookServer/Documents/

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

OpenDirectory Service

The entire Open Directory configuration can be saved with the archive feature.

Filetype	Location
Configuration files	/etc/openldap/slapd.conf
Data: (default locations)	/etc/openldap/ (stop slapd, and then backup with slapcat)

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 Dovecot

File type	Location
Configuration files	/etc/dovecot/dovecot.conf
	/etc/dovecot/partition_map.conf
Data: (default locations)	/var/mail/
	/var/spool/imap

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
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

Tomcat App Server

Filetype	Location
Configuration files	/Library/Tomcat/conf/
Data: (default locations)	/Library/Tomcat/webapps/

Web Service

File type	Location
Configuration files	/etc/apache2/* (for Apache 2.2)
	/etc/httpd/* (for Apache 1.3)
	/etc/webperfcache/*

File type	Location
	/Library/Keychains/System.keychain
Data: (default locations)	/Library/WebServer/Documents/ /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 computer and a file storage pool for 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.

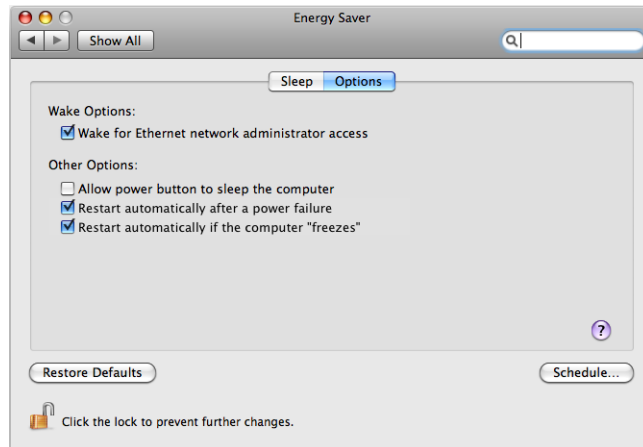
The following illustration is an example of an Xserve connected to a UPS via a network:



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

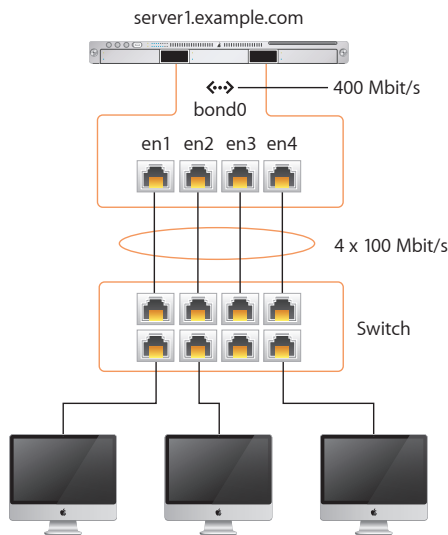
Although not common, the failure of a switch, cable, or network interface card can cause your server to become unavailable. To eliminate these single points of failure, you can use link aggregation or trunking. This technology, also known as IEEE 802.3ad, is built into Mac OS X and Mac OS X Server.

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 0xffffffff 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.

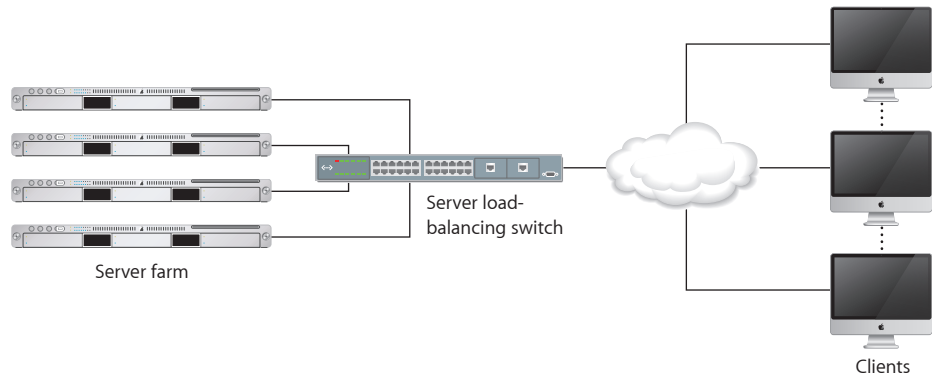
- 6 To view more information about a link, click the corresponding entry in the list.

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

Monitoring

8

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.

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?
- **Debugging:** Recurring problems can be analyzed and fixed if properly tracked. Even if you don't control any source code, good debugging logs and data can increase the ability of the developer to address your issues. How can you capture what is going wrong? How often? Does that data tell you what you need to know? Are they problems you can fix on your end, or do you need vendor support?

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 v10.6 activity from any computer with Snow Leopard. 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`, `launchd`, and `cron`.

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
É
```

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.

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.

- 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.

Server Status Notification in Server Admin

Server Admin has an easy to use notification system that can keep you informed of your server's hard disk status, software status and certificate 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 for the server.
- There is a certificate which has or will soon expire.

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.

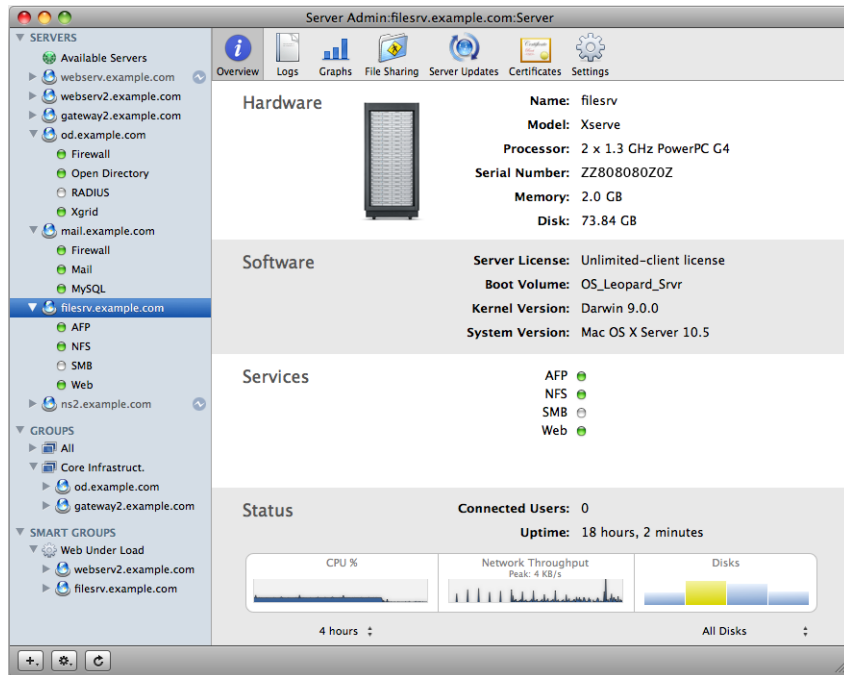
Monitoring Server Status Overviews Using Server Admin

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.

Remote Kernel Core Dumps

A kernel panic is a type of error that occurs when the core (kernel) of an operating system receives an instruction in an unexpected format, or that it fails to handle properly. A kernel panic may also follow when the operating system is not able to recover from a different type of error. A kernel panic can be caused by damaged or incompatible software or, more rarely, damaged or incompatible hardware.

When a server kernel panics it abruptly halts all normal system operations. Usually, a kernel process named `panic()` outputs an error message to the console, and stores debugging information in non-volatile memory to be written to a crash log file upon restarting the computer. Saving the memory contents of the core and associated debugging information is called a “core dump.”

This debugging information is highly technical, but system administrators can use this information to:

- record details about which machines are panicking and why.

For example, if you manage a large number of Mac OS X Servers, you might want to monitor which servers are panicking and why. You can use this information to determine how frequently kernel panics occurs, whether there are any common symptoms, and, most importantly, whether any third-party kernel extensions are involved.

- perform offline debugging on high-availability systems.

If you manage a high-availability server and you have problems with the server panicking, you can capture a kernel core dump, immediately restart the server, and then debug the problem without interrupting service.

For more information on debugging core dumps see Developer Technical Note #2118

<http://developer.apple.com/technotes/tn2004/tn2118.html#SECDEBUG>

You can configure a Mac OS X Server computer so that, when the machine panics, it transmits a core dump of the kernel to a remote core dump server via TCP/IP. The core dump server uses a daemon to collect the kernel core dump from the client and writes it to a file on the hard disk. You can then analyze the core dump using a variety of tools, most notably GDB.

CAUTION: The core dump of kernel memory is sent to the server in the clear. It's possible that this data might include sensitive information. You should therefore configure your network so that this data can't be seen by unauthorized persons (for example, use switched hubs, a firewall, or a VPN).

If you want to use a FireWire connection to transmit a core dump (a useful alternative when the kernel panic on the client system involves the built-in Ethernet driver, or some other network code), see the Read Me file in the FireWire SDK for Mac OS X describes the setup process for using FireWire to transmit a core dump.

The following sections contain information necessary to set up a remote listening server, which receives core dump information from panicked computers, and set up a server to send its core dump information to the remote listening server via TCP/IP over ethernet.

Setting Up a Core Dump Server

You can use any Mac OS X version 10.5 or later computer to be a core dump server which fits the following criteria. The core dump server must:

- have a static IP address.
- be IPv4 network-accessible to all the clients using UDP port 1069.
You cannot put the core dump server behind a firewall or NAT unless all the clients using it are also behind it. You cannot use IPv6-only addresses for the server.
- have enough disk storage space for multiple dumps.
In general, core dumps are fairly large. Core dumps can be as small as 200 MB to 500 MB but they can be much larger, depending on the kernel map size, physical memory size, memory usage during the panic, and other factors. Make sure you have a lot of free disk space.

If you need to set up a core dump server on a computer running a system earlier than Mac OS X version 10.5, more extensive configuration is needed. See Developer Technical Note #2118:

<http://developer.apple.com/technotes/tn2004/tn2118.html>

Setting up a core dump server:

- 1 Create a core dump directory named “PanicDumps,” owned by user “root,” and group “wheel,” which is writable by everyone.

Using the command line, type:

```
sudo mkdir /PanicDumps
sudo chown root:wheel /PanicDumps
sudo chmod 1777 /PanicDumps
```

- 2 Activate the core dump server process (kdumpd).

Using the command line, type:

```
sudo launchctl load -w /System/Library/LaunchDaemons/com.apple.kdumpd.plist
```

Once this command is executed, the core dump server process will always automatically start. This step does not have to be repeated when the server restarts.

- 3 Verify that the core dump server process is running.

Using the command line, type:

```
sudo launchctl list | grep kdump
```

The result should list com.apple.kdumpd.

- 4 Make sure UDP port 1069 is open for core dump connections.

Now that the core dump server is active, you need to configure Mac OS X computers as clients to send their kernel panic information to this server. See “Setting Up a Core Dump Client.”

Setting Up a Core Dump Client

A core dump client sends its kernel panic debug information to the core dump server address specified in its NVRAM settings. The information is transmitted at the time of the panic, so before restarting the computer, allow some time after the panic for the data to be sent to the server. The time necessary depends on the file size of the core dump and the speed of the network connection between the client and server.

If you need to make client 10.5 or earlier, see [<http://developer.apple.com/technotes/tn2004/tn2118.html>]

Setting up a core dump client:

- 1 Modify the `boot-args` NVRAM variable to include the `debug` flag `0x0400`, and the `_panicd_ip` flag with the IP address of the core dump server.

The following example uses the core dump server IP address as 192.168.1.250. Substitute the IP address of your own core dump server.

```
sudo nvram boot-args="debug=0x0400 _panicd_ip=192.168.1.250"
```

Important: The `boot-args` NVRAM variable may be reset whenever you install new system software, including software updates, and whenever you change the startup disk using System Preferences.

- 2 If the core dump client is running Mac OS X Server, modify the `watchdogtimerd` behavior to either keep it from restarting the server before the core dump is complete, or modify the amount of time it waits before restarting the server.

To disable automatic restarting, turn off the “Restart automatically after a power failure” option in the Options tab of the Energy Saver System Preferences pane.

To increase the amount of time before automatic restarting, add a “count” program argument larger than 6 (but smaller than 480) to the `watchdogtimerd` configuration file at `/System/Library/LaunchDaemons/com.apple.watchdogtimerd.plist`. See the `watchdogtimerd(8)` man page for more information about the arguments and options.

- 3 Restart the computer for the settings to take effect.

For additional NVRAM debug flags which are useful in core dump debugging, see Developer Technical Note #2118, sub-section “Debug Flags in Depth.”

<http://developer.apple.com/technotes/tn2004/tn2118.html>

Configuring Common Core Dump Options

By default, core dumps happen using UDP port 1069 over the built-in Ethernet (en0) interface, and the resulting files are stored in /PanicDumps on the core dump server. However, you can configure the core dump to use:

- an alternate UDP port
- an alternate network interface
- an alternate file destination
- a specific network router

Changing any of these options requires that you restart the computers to reload the new settings. All settings assume the core dump client and the core dump server are using Mac OS X 10.5 or later.

Option	Action
To set an alternate UDP port...	<p>On the core dump server, change the SockServiceName string property from 1069 to the desired port in /System/Library/LaunchDaemon/com.apple.kdump.plist.</p> <p>On the core dump client, add the <code>_paniced_port</code> flag to the NVRAM <code>boot-args</code>. For example, to change it to UDP port 12345, add "<code>_paniced_port=12345</code>" to the list of <code>boot-args</code> flags.</p>
To set an alternate network interface...	<p>On the core dump client, add the <code>kdp_match_name</code> flag to the NVRAM <code>boot-args</code>. For example, to change it to always use en1, add "<code>kdp_match_name=en1</code>" to the list of <code>boot-args</code> flags after the <code>_paniced_ip</code> flag.</p> <p>AirPort interfaces cannot be used to transmit core dumps.</p>
To set an alternate file destination...	<p>On the core dump server, change the expected directory location in the /System/Library/LaunchDaemons/com.apple.kdumpd.plist file, ProgramArguments string, then reload the kdumpd process.</p>
To specify a network router...	<p>On the core dump client, add the <code>_router_ip</code> flag to the NVRAM <code>boot-args</code>. For example, to change it to use the router 10.0.0.1, add "<code>_router_ip=10.0.0.1</code>" to the list of <code>boot-args</code> flags after the <code>_paniced_ip</code> flag.</p>

If you want to change the location of the core dump directory, you must change the expected directory location in the com.apple.kdumpd.plist file, then reload the process.

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 (TBD update)

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 183.

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:

- Select File: 1 (`snmpd.conf`)
- Select section: 5 (System Information Setup)
- Select section: 1 (The [typically physical] location of the system.)
- The location of the system: type text string here — such as “`server_room`”
- Select section: f (finish)
- Select section: f (finish)
- 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.

Para

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.

Additional Information about SNMP

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

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.

•

-

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 v10.6. 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.

The logs are set to roll (compress and rename the log file) every 5MB.

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 <code>syslog.conf</code>	Amount of detail
None	<code>.none</code>	None
Emergency	<code>.emerg</code>	Least
Alert	<code>.alert</code>	
Error	<code>.err</code>	
Warning	<code>.warn</code>	
Notice	<code>.notice</code>	

Level name	Level indicator in syslog.conf	Amount of detail
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 -enableslapdlog
```

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 `syslogd.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.

Provide increased server responsiveness to clients and reduce server load with Push Notification Server.

Mac OS X Server v10.6 uses an XMPP Pubsub architecture for the Push Notification Server. XMPP Pubsub is an open standard extension to XMPP (XEP-060) that allows servers and clients to communicate on an “as needed” basis rather than clients continually asking the server for updates.

A service (like iCal or mail) maintains a simple connection with the client and the service informs the client that there is new data. This differs from previous methods, where calendar or mail clients contacted the server at regular intervals, requesting new data, if present. With the previous method of notification, the server has to attend to each client, regardless of whether the client has new data waiting for it. By using the new push method of client updating, only the clients with new data are ever contacted, and even then only as needed.

The push method yields immediate benefits to server responsiveness to clients and reduces the overall workload of the server.

Starting and Stopping Push Notification

When you start push notification on a server, the service broadcasts its availability on the local network to any other services which support it. This means that when a different server turns on a service that supports push notification, the push notification server address automatically populates in the settings of the pushing service. You must still enable Push Notification support for the pushing service before it works.

Additionally, you can choose to encrypt the data passed between the client and the push server by choosing an SSL certificate. This does not encrypt the data between the client and the pushing service. To encrypt transport between the pushing service and the client, you still need to enable SSL with the pushing service.

To enable Push Notification:

- 1 Use Server Admin to connect to the desired server.
- 2 Enable administration of Push Notification.

This only needs to be done when you administer is for the first time. For more information about adding a service to the administered services list, see “Adding and Removing Services in Server Admin.”

- 3 Select Push Notification from the list of administered services for the desired server.
- 4 To enable SSL transport encryption for the notifications, select the desired SSL certificate.
- 5 Click the “Start Push Notification” or “Stop Push Notification” as needed.

Index

A

access

- ACLs 55, 75
- IP address restrictions 52
- Keychain Access Utility 66
- LDAP 21
- remote installation 83
- SACLs 75
- user 134, 149
- See also* permissions

accounts. *See* Workgroup Manager; user accounts

ACLs (access control lists) 55, 75

addresses. *See* IP addresses

Administer permission level 151

administrator 74, 75, 76, 151, 152

administrator computer 82, 126, 127

AFP (Apple Filing Protocol) service 22, 189

Apple Remote Desktop (ARD) 49, 133, 190

archiving server data 32, 36

asr tool 36, 87

authentication

- and TLS[authentication
TLS] 54
- Kerberos 21, 57, 58, 113
- key-based SSH 72, 73
- keychain services 157
- MS-CHAPv2 111
- Open Directory 57
- overview 56
- passwords 57, 78, 98
- RADIUS 21, 58, 157
- SASL 57
- Server Admin 38, 39, 63
- single sign-on 58
- standalone server 112
- users 56, 58, 73, 111
- Workgroup Manager 153
- See also* certificates

authorization 56

- See also* authentication

B

backups

- command-line tools 36
- critical files 157
- media types 35
- policy considerations 31, 35
- rotation scheme 34
- scheduling 33
- server setup data 119, 121
- Time Machine 36
- types 33
- validation of 35

Berkeley Software Distribution. *See* BSD

BSD (Berkeley Software Distribution) 23

C

calendar service. *See* iCal service

Certificate Authority (CA)

- creating 66
- creating certificates from 68
- distributing to clients 70
- introduction 60
- overview 61
- requesting certificates from 63, 64, 65, 68
- See also* PKI

Certificate Manager 63, 68

Certificate Signing Request. *See* CSR

certificates

- and Server Admin 63
 - Server Admin 150
- and services 72
- creating 65, 66, 68
- deleting 71
- editing 70
- identities 61
- importing 68
- managing 69
- overview 59, 60
- preparing 64
- private keys 60
- public keys 60
- renewing 71
- requesting 64, 65
- root 66
- self-signed 61, 65

- client computers and NetBoot 27
 - clients
 - certificates 70
 - and NetBoot 27
 - client-side logging 189
 - group accounts 155
 - See also* users
 - command-line tools
 - backup tools 36
 - and permissions
 - permissions 152
 - daemon control 171
 - disk space monitoring 176
 - erasing disks 97
 - installing server software 103
 - partitioning disks 94
 - restoration tools 36
 - server administration 48
 - computer lists 153, 154, 155
 - computer name 134
 - computers, administrator 82, 126, 127
 - computer-to-computer network 166
 - computer-to-switch network 167
 - computer-to-switch-pair network 167
 - concatenated RAID set 95
 - configuration
 - advanced 112
 - authentication 57
 - automatic 118, 119, 120, 121
 - batch setup for multiple servers 116
 - connecting to network 109, 166, 167
 - DHCP 81, 112
 - directory connection 111, 112
 - Ethernet 109
 - interactive 113, 115, 116
 - introduction 18, 108
 - link aggregation 168
 - logs 123
 - Open Directory 110, 111, 112, 124
 - postponing 108
 - remote server 113, 115, 116
 - saving setup data 119, 120, 121
 - server infrastructure 29
 - server types 18
 - services 124, 125, 157
 - settings overview 109
 - SSL 150
 - standalone server 110
 - status checking 122, 123
 - troubleshooting 123
 - types of 108
 - Console 175
 - CSR (Certificate Signing Request) 63, 64, 65, 68
- D**
- daemons, overview 171
 - Darwin (core operating system) 23
 - Date & Time preferences 134
 - df tool 176
 - DHCP (Dynamic Host Configuration Protocol)
 - service 29, 81, 112
 - digital signature 150
 - directories. *See* folders; domains, directory; directory services
 - directory services
 - advanced configuration 112
 - directory domains 20, 111, 112, 155
 - logs 188
 - planning of 26
 - See also* Open Directory
 - Directory, overview 46
 - disk images
 - encrypting 56
 - installing with 27, 47, 86, 89
 - Disk Utility 56, 93, 95, 97
 - disks
 - command-line management of 94, 97, 176
 - erasing free space 97
 - installation preparation 90, 92, 93, 94, 95, 96, 97
 - mirroring 95
 - monitoring tools 175
 - partitions 86, 92, 93, 94, 95, 97
 - quotas 27
 - See also* RAID
 - diskspacemonitor tool 176
 - diskutil tool 94, 96, 97
 - ditto tool 36
 - DMZ, network 52
 - DNS (Domain Name System) service 29, 81
 - documentation 12, 13, 14
 - Domain Name System. *See* DNS
 - domains, directory 20, 111, 112, 155
 - drives. *See* disks
 - du tool 176
 - DVDs, installation 85
 - Dynamic Host Configuration Protocol. *See* DHCP
- E**
- email. *See* mail service
 - emond daemon 186
 - encryption 54, 55, 56, 60, 121
 - See also* SSL
 - Ethereal packet sniffing tool 177
 - Ethernet 53, 109, 168
 - exporting service settings 149
- F**
- file services 22, 189
 - file sharing 150
 - file systems
 - backing up 36
 - choosing 90

See also volumes
File Transfer Protocol. *See* FTP
files
 backup 31, 35, 157
 configuration 188
 full file-level copies 33
 security 55, 56
 setup data 119, 120, 121
 shared secret 60
 storage considerations 27
FileVault 56
Firewall service 52, 53, 81, 159
folders 27, 55, 133
FTP (File Transfer Protocol) service 22
full file-level copies 33
full image backup type 33

G

Gateway Setup Assistant 157
group accounts 155
groups 131, 149, 151, 153, 154
Growl application 190

H

hardware requirements 16, 30, 80, 81, 95
help, using 12
HFS+J volume 91
HFSX volume 91
historical data collection 173
home folders 27, 133
host name
 local 134

I

iCal service 158
iChat service 158
identity certificates. *See* certificates
images. *See* NetInstall; disk NetBoot; disk images
importing
 certificates 68
 service settings 149
incremental backups 33
infrastructure requirements 29
Inspector 155
installation
 administrator computer 82
 collecting information 81
 command-line method 103
 disk preparation 90, 92, 94, 95, 96, 97
 from earlier OS versions 25, 28, 79, 82
 identifying servers 98
 infrastructure requirements 29
 with disk images 27, 47, 86, 89
 integration strategy 28
 interactive 99, 101, 102

multiple server 105
network services setup 81
overview 79
planning for 24, 25, 26, 27, 28
postponing setup after 108
remote access 82, 83, 98, 101
server installation disc 82
server software 82, 103
starting up for 83, 85, 86, 89
system requirements 80, 81
updating 107
upgrading 106
installer tool 103, 106
IP addresses
 access restriction 52
 changing server 31, 146
 and firewalls 81
 overview 22
 remote server installation 84, 98
 servers on different subnets 114
IPv6 addressing 22

J

journaling, file system 91

K

KDC (Kerberos Key Distribution Center). *See* Kerberos
Kerberos 21, 57, 58, 113
key-based authentication 72, 73
Keychain Access Utility 66
keychain services 157

L

LACP (Link Aggregation Control Protocol) 165
launchctl tool 172
launchd daemon 36, 37, 171
LDAP (Lightweight Directory Access Protocol)
 service 21
LDAPv3 servers 58
link aggregation 165, 166, 167, 168, 169
Link Aggregation Control Protocol. *See* LACP
load balancing 170
local directory domain, standalone server 112
login, authenticating 72, 73
logs
 monitoring 175, 186, 187, 188, 189
 troubleshooting setup 123
 web services 161

M

MAC (media access control) addresses 53, 98
Mac OS X
 administration from 127
 installation considerations 82
 upgrading from 106

- Mac OS X Server
 - configuration 111
 - integration strategy 28
 - introduction 16, 18
 - and UNIX 23
 - supported standards 20
 - system requirements 16
 - See also* configuration; installation
- mail service 21, 157, 159
- managed preferences, defining 155
- media, streaming. *See* streaming media
- migration 25, 27, 28
- mirroring, disk 95
- mobile accounts 133
- Monitor permission level 151
- MS-CHAPv2 authentication 111
- MySQL service 160

N

- Nagios application 190
- NAT (Network Address Translation) 159
- NetBoot service 27, 47, 89
- NetInstall 47, 89
- Network Address Translation. *See* NAT
- Network File System. *See* NFS
- network interfaces 134
- network services
 - DHCP 29, 81, 112
 - DNS 29, 81
 - installation 81
 - NAT 159
 - NTP 133, 134
 - planning for 29
 - VLAN 53
 - See also* IP addresses
- network time protocol. *See* NTP
- networks
 - connection configurations 109, 166, 167
 - environment for installation 79
 - Ethernet 53, 109, 168
 - initial server setup connection 109
 - monitoring tools 177, 182, 183, 184, 185
 - security 52, 53, 54, 55
- NFS (Network File System) 22
- notification system 44, 134, 158, 177, 186
- NTP (network time protocol) 133, 134

O

- Open Directory
 - authentication 57
 - logs 189
 - and SACLs 75
 - setup 110, 111, 112, 124
- Open Directory master 81
- Open Directory replica 57, 164
- open source modules

- Kerberos 21, 57, 58, 113
- OpenLDAP 21
- OpenSSL 54
- PHP 160
- See also* Open Directory
- OpenLDAP 21
- OpenSSL 54
- operating environment requirements 164

P

- PackageMaker 47
- packets, data, filtering of 52
- partitions, disk 86, 92, 94, 95, 97
- passwords 57, 78, 98
- permissions
 - administrator 74, 75, 151, 152
 - files 55
 - folder 55
 - SACL 75
 - types 55
- php configuration files 160
- physical infrastructure requirements 29
- PKI (public key infrastructure) 54, 59, 60
- portable computers 133
- Portable Operating System Interface. *See* POSIX
- ports
 - Ethernet 109
 - list of 129
 - status of 129
 - TCP 72
- POSIX (Portable Operating System Interface) 55
- preferences 155
- presets 154
- private key 60, 61
- privileges, administrator 75, 151, 152
- See also* permissions
- PropertyListEditor 47
- protocols
 - file service 22, 189
 - network service 29, 81, 112, 133, 146
 - overview 22
 - See also* specific protocols
- public key certificates. *See* certificates
- public key cryptography 72
- public key infrastructure. *See* PKI

Q

- QuickTime Streaming Server (QTSS) 47, 158
- quotas, disk space 27

R

- RADIUS (Remote Authentication Dial-In User Service) 21, 58, 157
- RAID Admin 175

- RAID (Redundant Array of Independent Disks) 24, 27, 94, 95, 96
 - real-time monitoring 173
 - Remote Authentication Dial-In User Service. *See* RADIUS
 - remote servers
 - accessing 83
 - Apple Remote Desktop 49, 133, 190
 - configuration 113, 115, 116
 - identifying 98
 - installing from or to 82, 83, 98, 101
 - replication 57, 164
 - requirements
 - hardware 16, 30, 80, 81, 95
 - infrastructure 29
 - operating environment 164
 - software 80, 82
 - restart, automatic 163
 - restoration, data 31, 34
 - root certificate 66
 - rsync tool 36
- S**
- SACLs (service access control lists) 75
 - SASL (Simple Authentication and Security Layer) 57
 - Secure Empty Trash 56
 - secure SHell. *See* SSH
 - Secure Sockets Layer. *See* SSL
 - Secure VM 56
 - security
 - administrator 74, 75
 - authorization 56
 - best practices 76
 - file 55, 56
 - Firewall service 52, 53, 81, 159
 - installation 82
 - network 52, 53, 54, 55
 - overview 51
 - physical 51
 - SASL 57
 - service level 75
 - settings 150
 - SSH 72, 73, 83, 133, 157
 - SSL 54, 59, 60, 63, 150
 - TLS 54
 - See also* access; authentication; certificates; SSL
 - self-signed certificates 61, 65
 - serial number, server 85, 98
 - Server Admin 129
 - access control 149
 - as administration tool 130
 - authentication 38, 39, 63
 - certificates 63, 150
 - customizing 40
 - notification system 177
 - opening 38, 39, 63
 - overview 11, 18, 38, 39, 63
 - server status 178
 - and system imaging 47
 - service management 148
 - Server Assistant 41, 101, 108, 113
 - Server Message Block protocol. *See* SMB
 - Server Monitor 44, 174
 - Server Status Dashboard widget 174
 - servers
 - adding 130
 - administration tools 38, 48, 49, 126, 127
 - basic settings 109, 132
 - groups of 131
 - infrastructure requirements 29
 - load balancing 170
 - reliability tools 161, 162, 163, 164, 165, 166, 168, 169
 - relocation considerations 31
 - removing 130
 - serial numbers for 85, 98
 - standalone 110, 112
 - startup 83, 89
 - status monitoring 173, 174, 175, 177, 178
 - time 133, 134
 - troubleshooting 123
 - service access control lists. *See* SAACLs
 - services
 - access control 134, 149
 - exporting settings 149
 - importing settings 149
 - management of 157
 - planning for distribution of 26
 - security 72, 75
 - setup 124, 125, 157
 - viewing 134, 147
 - See also* specific services
 - setup procedures. *See* installation; configuration
 - share points 55, 150
 - shared directory domain 21, 110
 - shared secret files 60
 - Simple Network Management Protocol. *See* SNMP
 - single points of failure 161
 - single sign-on authentication 58
 - slapd daemon 189
 - SMB (Server Message Block) protocol 22
 - snapshots, data 33
 - SNMP (Simple Network Management Protocol)
 - definition 22
 - settings 133
 - as monitoring tool
 - monitoring 182, 183, 184, 185
 - snmpd daemon 183
 - Software Update service 107
 - srm UNIX utility 56
 - SSH (secure SHell host) 72, 73, 83, 133, 157
 - SSL (Secure Sockets Layer) 54, 59, 60, 63, 150

- standalone server 110, 112
- standard configuration type 18
- streaming media 27, 47, 158
- striping 94
- subnets 109, 114
- syslog configuration file 187
- syslogd daemon 186
- System Image Utility 47

T

- TCP (Transmission Control Protocol) 52, 72
- tcpdump tool 177
- time server 133, 134
- TLS (Transport Layer Security) protocol 54
- Transmission Control Protocol. *See* TCP
- Transport Layer Security protocol. *See* TLS
- troubleshooting server operation 123

U

- UDP (User Datagram Protocol) 52
- UNIX 23
- upgrading
 - from Mac OS X 106
 - from previous server versions 25, 28
 - and saved setup data 120
 - vs. migration 25, 28
- UPS (uninterruptible power supply) 162, 163
- user accounts
 - authentication 58
 - group 155
 - managed preferences 155
 - management of 153, 154
 - mobile 133
 - passwords 57
 - setup 125
 - See also* users
- User Datagram Protocol. *See* UDP
- users
 - access control 134, 149
 - administrative access for 74, 75
 - authentication 56, 58, 73, 111
 - certificates 61
 - disk space quotas 27
 - groups 149, 151, 153, 154
 - home folders 27, 133
 - management of 153, 154
 - permissions 151
 - Windows 27
 - See also* clients; user accounts; Workgroup Manager

V

- Virtual Private Network. *See* VPN
- VLAN (virtual local area network) 53
- VNC (virtual network computing) 80, 83, 102, 105

- volumes
 - backing up 36
 - erasing 97
 - RAID 94, 95
 - startup 83, 89
 - supported 91
 - and partitioning 92, 94

W

- web services 160
- web technologies 22
- weblog service 161
- wikis 161
- Windows NT 28
- Windows users 27
- Workgroup Manager 152
 - administering accounts 153
 - administration overview 153
 - authentication 153
 - customizing 44, 156
 - opening 42, 153
 - overview 42, 43

X

- Xgrid 2 service 157
- Xgrid Admin 49
- Xserve
 - hardware installation instructions 81
 - VLAN support 53
 - and Server Monitor 44
 - and server reliability
 - server reliability 162, 163