




Mac OS X Server

Network Services Administration
For Version 10.4 or Later

 Apple Computer, Inc.
© 2005 Apple Computer, Inc. All rights reserved.

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

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

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

Use of the "keyboard" Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

Apple, the Apple logo, AirPort, AppleScript, AppleShare, AppleTalk, Mac, Mac OS, Macintosh, Power Mac, Power Macintosh, QuickTime, Sherlock, and WebObjects are trademarks of Apple Computer, Inc., registered in the U.S. and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

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

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

019-0165/3-24-05

Contents

Preface	9 About This Guide
	9 What's New in Version 10.4?
	9 What's in This Guide
	10 Using This Guide
	10 Using Onscreen Help
	11 The Mac OS X Server Suite
	12 Getting Documentation Updates
	12 Getting Additional Information
Chapter 1	15 Linking Your Network to the Internet
	15 Understanding the Gateway Setup Assistant
	16 Using the Gateway Setup Assistant
	16 Sample Configurations
	17 Connecting a Wired LAN to the Internet
	18 Connecting a Wired LAN And Wireless Clients To The Internet
	19 Connecting a Wireless LAN To The Internet
Chapter 2	23 DHCP Service
	23 Before You Set Up DHCP Service
	24 Creating Subnets
	24 Assigning IP Addresses Dynamically
	24 Using Static IP Addresses
	25 Locating the DHCP Server
	25 Interacting With Other DHCP Servers
	25 Using Multiple DHCP Servers on a Network
	25 Assigning Reserved IP Addresses
	25 Getting More Information on the DHCP Process
	26 Setting Up DHCP Service for the First Time
	26 Managing DHCP Service
	26 Starting and Stopping DHCP Service
	27 Creating Subnets in DHCP Service
	27 Changing Subnet Settings in DHCP Service
	28 Deleting Subnets From DHCP Service

28	Disabling Subnets Temporarily
28	Changing IP Address Lease Times for a Subnet
29	Setting the DNS Server for a DHCP Subnet
29	Setting LDAP Options for a Subnet
30	Setting WINS Options for a Subnet
30	Assigning Static IP Addresses Using DHCP
31	Removing or Changing Static Address Maps
31	Monitoring DHCP Service
31	Viewing the DHCP Status Overview
32	Setting the Log Detail Level for DHCP Service
32	Viewing DHCP Log Entries
32	Viewing the DHCP Client List
33	Common Network Configurations That Use DHCP
36	Where to Find More Information

Chapter 3

37	DNS Service
38	Before You Set Up DNS Service
38	DNS and BIND
38	Setting Up Multiple Name Servers
38	Setting Up DNS Service for the First Time
41	Managing DNS Service
41	Starting and Stopping DNS Service
41	Enabling or Disabling Zone Transfers
42	Enabling or Disabling Recursion
42	Managing DNS Zones
43	Adding a Primary Zone
44	Adding a Secondary Zone
44	Duplicating a Zone
45	Modifying a Zone
45	Deleting a Zone
45	Using an Existing Zone File
46	Managing DNS Machine Records
47	Adding a Machine Record to a DNS Zone
48	Modifying a Machine Record in a DNS Zone
48	Deleting a Machine Record From a DNS Zone
49	Monitoring DNS
49	Viewing DNS Service Status
49	Viewing DNS Log Entries
49	Changing DNS Log Detail Levels
50	Changing DNS Log File Location
50	Securing the DNS Server
50	DNS Spoofing
51	Server Mining

52	DNS Service Profiling
52	Denial of Service (DoS)
52	Service Piggybacking
53	Common Network Administration Tasks That Use DNS Service
53	Setting Up MX Records
55	Setting Up Namespace Behind a NAT Gateway
56	Network Load Distribution (aka Round Robin)
56	Setting Up a Private TCP/IP Network
57	Hosting Several Internet Services With a Single IP Address
58	Hosting Multiple Domains on the Same Server
58	Where to Find More Information

Chapter 4

59	IP Firewall Service
60	Basic Firewall Practices
61	Firewall Startup
62	Understanding Firewall Rules
62	What is a Firewall Rule?
64	Using Address Ranges
64	Rule Mechanism and Precedence
64	Multiple IP Addresses
65	Setting Up Firewall Service for the First Time
66	Managing Firewall Service
66	Managing Panther Server 10.3 Firewalls with Tiger Server 10.4 Server Admin
66	Starting and Stopping Firewall Service
67	Creating an Address Group
67	Editing or Deleting an Address Group
68	Duplicating an Address Group
68	Opening the Firewall for Standard Services
69	Adding to the Services List
70	Editing or Deleting Items in the Services List
70	Creating an Advanced IP Firewall Rule
71	Editing or Deleting Advanced IP Firewall Rules
72	Changing the Advanced IP Firewall Rule Order
72	Enabling Stealth Mode
72	Resetting an Unreachable Server
73	Monitoring Firewall Service
73	Understanding the Active Rules Panel
73	Viewing the Firewall Status Overview
74	Viewing the Active Firewall Rules
74	Setting Up Logs for Firewall Service
74	Viewing the Firewall Log
75	Viewing Denied Packets
76	Viewing Packets Logged by Firewall Rules

76	Troubleshooting Advanced IP Firewall Rules
77	Practical Examples
77	Using IP Firewall with NAT
77	Block Web Access to Internet Users
78	Logging Internet Access by Local Network Users
79	Block Junk Mail
79	Allow a Customer to Access the Apple File Server
80	Common Network Administration Tasks That Use Firewall Service
80	Preventing Denial of Service (DoS) Attacks
81	Controlling or Enabling Peer-to-Peer Network Usage
81	Controlling or Enabling Network Game Usage
82	Port Reference
86	Where to Find More Information

Chapter 5

87	NAT Service
87	Using NAT with Other Network Services
88	NAT LAN Configuration Overview
89	Starting and Stopping NAT Service
89	Configuring NAT Service
90	Creating a Gateway Without NAT
90	Configuring Port Forwarding
91	Port Forwarding Examples
93	Monitoring NAT Service
93	Viewing the NAT Status Overview
93	Common Network Administration Tasks That Use NAT
93	Linking a LAN to the Internet Through One IP Address
95	Setting Up LAN Party For Gaming
95	Setting Up “Virtual Servers”
97	Where to Find More Information

Chapter 6

99	VPN Service
99	VPN and Security
99	Transport Protocols
100	Authentication Method
101	Before You Set Up VPN Service
101	Configuring other Network Services for VPN
102	Managing VPN Service
102	Starting or Stopping VPN Service
102	Enabling and Configuring L2TP Transport Protocol
102	Enabling and Configuring PPTP Transport Protocol
103	Configuring Additional Network Settings for VPN Clients
103	Configuring VPN Network Routing Definitions
105	Limiting VPN Access to Certain Users or Groups

	106	Limiting VPN Access to Certain Incoming IP Addresses
	107	Additional Configuration Instructions
	109	Monitoring VPN Service
	109	Viewing a VPN Status Overview
	109	Setting the Log Detail Level for VPN Service
	109	Viewing the VPN Log
	110	Viewing VPN Client Connections
	110	Common Network Administration Tasks That Use VPN
	110	Linking a Computer at Home With a Remote Network
	112	Accessing a Single Computing Asset Behind a Remote Network Firewall
	112	Linking Two or More Remote Network Sites
	116	Where to Find More Information
Chapter 7	117	NTP Service
	117	How NTP Works
	118	Using NTP on Your Network
	118	Setting Up NTP Service
	119	Configuring NTP on Clients
	119	Where to Find More Information
Chapter 8	121	VLAN Support
	121	Understanding VLANs
	121	Setting Up Client Membership to a VLAN
	122	Where to find more information
Chapter 9	123	IPv6 Support
	124	IPv6 Enabled Services
	124	IPv6 Addresses in the Server Admin
	124	IPv6 Addresses
	124	Notation
	125	IPv6 Reserved Addresses
	125	IPv6 Addressing Model
	125	IPv6 Address Types
	126	Where to Find More Information
Glossary	127	
Index	139	

About This Guide

This guide explains how to configure and administer Mac OS X Server Network Services.

What's New in Version 10.4?

Mac OS X Server version 10.4 has many improvements over version 10.3 and additional features. Among them are:

- New Gateway Setup Assistant
- Revised and improved DNS interface
- Static IP address mapping through DHCP
- Revised and improved Firewall interface
- Expanded VPN help
- Expanded NAT help
- VLAN support information

What's in This Guide

This guide is divided into nine chapters and a glossary:

- Chapter 1, "Linking Your Network to the Internet," on page 15, tells you how to use Gateway Setup Assistant to link your network to the Internet.
- Chapter 2, "DHCP Service," on page 23, tells you how to configure and use DHCP to assign IP addresses on your network.
- Chapter 3, "DNS Service," on page 37, tells you how to use Mac OS X Server as a domain name server.
- Chapter 4, "IP Firewall Service," on page 59, tells you how to maintain network security using a firewall.
- Chapter 5, "NAT Service," on page 87, tells you how to configure and use NAT to connect many computers to the Internet with only one public IP address.
- Chapter 6, "VPN Service," on page 99, tells you how to configure and use VPN to allow remote users to access your private LAN securely.
- Chapter 7, "NTP Service," on page 117, tells you how to enable your server as a time server.

- Chapter 8, “VLAN Support,” on page 121, tells you about VLAN support for some server hardware configurations.
- Chapter 9, “IPv6 Support,” on page 123, tells you about IPv6 and which services support IPv6 addressing.
- “Glossary” on page 127, contains definitions of the terms used in this guide.

Using This Guide

Each chapter covers a specific network service. Read any chapter that’s about a service you plan to provide to your users. Learn how the service works, what it can do for you, strategies for using it, how to set it up for the first time, and how to administer it over time.

Also take a look at chapters that describe services with which you’re unfamiliar. You may find that some of the services you haven’t used before can help you run your network more efficiently and improve performance for your users.

Most chapters end with a section called “Where to Find More Information.” This section points you to websites and other reference material containing more information about the service.

Using Onscreen Help

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

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

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

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

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

The Mac OS X Server Suite

The Mac OS X Server documentation includes a suite of guides that explain the services and provide instructions for configuring, managing, and troubleshooting the services.

All of the guides are available in PDF format from:

www.apple.com/server/documentation/

This guide...	tells you how to:
<i>Mac OS X Server Getting Started for Version 10.4 or Later</i>	Install Mac OS X Server and set it up for the first time.
<i>Mac OS X Server Upgrading and Migrating to Version 10.4 or Later</i>	Use data and service settings that are currently being used on earlier versions of the server.
<i>Mac OS X Server User Management for Version 10.4 or Later</i>	Create and manage users, groups, and computer lists. Set up managed preferences for Mac OS X clients.
<i>Mac OS X Server File Services Administration for Version 10.4 or Later</i>	Share selected server volumes or folders among server clients using these protocols: AFP, NFS, FTP, and SMB/CIFS.
<i>Mac OS X Server Print Service Administration for Version 10.4 or Later</i>	Host shared printers and manage their associated queues and print jobs.
<i>Mac OS X Server System Image and Software Update Administration for Version 10.4 or Later</i>	Use NetBoot and Network Install to create disk images from which Macintosh computers can start up over the network. Set up a software update server for updating client computers over the network.
<i>Mac OS X Server Mail Service Administration for Version 10.4 or Later</i>	Set up, configure, and administer mail services on the server.
<i>Mac OS X Server Web Technologies Administration for Version 10.4 or Later</i>	Set up and manage a web server, including WebDAV, WebMail, and web modules.
<i>Mac OS X Server Network Services Administration for Version 10.4 or Later</i>	Set up, configure, and administer DHCP, DNS, VPN, NTP, IP firewall, and NAT services on the server.
<i>Mac OS X Server Open Directory Administration for Version 10.4 or Later</i>	Manage directory and authentication services.
<i>Mac OS X Server QuickTime Streaming Server Administration for Version 10.4 or Later</i>	Set up and manage QuickTime streaming services.
<i>Mac OS X Server Windows Services Administration for Version 10.4 or Later</i>	Set up and manage services including PDC, BDC, file, and print for Windows computer users.
<i>Mac OS X Server Migrating from Windows NT for Version 10.4 or Later</i>	Move accounts, shared folders, and services from Windows NT servers to Mac OS X Server.

This guide...	tells you how to:
<i>Mac OS X Server Java Application Server Administration For Version 10.4 or Later</i>	Configure and administer a JBoss application server on Mac OS X Server.
<i>Mac OS X Server Command-Line Administration for Version 10.4 or Later</i>	Use commands and configuration files to perform server administration tasks in a UNIX command shell.
<i>Mac OS X Server Collaboration Services Administration for Version 10.4 or Later</i>	Set up and manage weblog, chat, and other services that facilitate interactions among users.
<i>Mac OS X Server High Availability Administration for Version 10.4 or Later</i>	Manage IP failover, link aggregation, load balancing, and other hardware and software configurations to ensure high availability of Mac OS X Server services.
<i>Mac OS X Server Xgrid Administration for Version 10.4 or Later</i>	Manage computational Xserve clusters using the Xgrid application.
<i>Mac OS X Server Glossary: Includes Terminology for Mac OS X Server, Xserve, Xserve RAID, and Xsan</i>	Interpret terms used for server and storage products.

Getting Documentation Updates

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

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

Getting Additional Information

For more information, consult these resources:

Read Me documents—important updates and special information. Look for them on the server discs.

Mac OS X Server website—gateway to extensive product and technology information. www.apple.com/macosx/server/

AppleCare Service & Support—access to hundreds of articles from Apple’s support organization. www.apple.com/support/

Apple customer training—instructor-led and self-paced courses for honing your server administration skills.
train.apple.com

Apple discussion groups—a way to share questions, knowledge, and advice with other administrators.
discussions.info.apple.com

Apple mailing list directory—subscribe to mailing lists so you can communicate with other administrators using email.
www.lists.apple.com

Linking Your Network to the Internet

1

You use Gateway Setup Assistant to link your network to the Internet. It guides you through initial setup of a server that will serve as a gateway between your private network and the Internet.

Understanding the Gateway Setup Assistant

The Gateway Setup Assistant helps you to quickly and easily set up a Mac OS X Server 10.4 to share your connection with the Internet with your local network. After making a few configuration choices, the assistant saves all the appropriate settings to start sharing the server's connection.

Depending on your configuration choices, the assistant will do the following when it is finished:

- Assign the server a static IP address for each internal network interface.
The address assigned is 192.168.x.1. The number used for x is determined by the network interface's order in the Network System Preference pane. For example, for the first interface, x is 0, for the second interface on the list, x is 1.
- Enable the DHCP to allocate addresses on the internal network, removing existing DHCP subnets.
- Set aside certain internal (192.168.x.x) addresses for DHCP use.
Without VPN enabled, each interface can allocate 192.168.x.2–192.168.x.254.
- Enable VPN (optional), to allow authorized external clients to connect to the local network.
VPN L2TP is enabled, so you must enter a shared secret for client connections to use.
- Set aside certain internal (192.168.x.x) addresses for VPN use.
If VPN is selected, half of the allotted IP addresses in the DHCP range are reserved for VPN connections. The addresses 192.168.x.128–192.168.x.254 are allotted to VPN connections.

- Enable the IP Firewall to help secure the internal network. Address groups are added for each internal network interface, with all traffic allowed from the newly created DHCP address ranges to any destination address.
- Enable NAT on the internal network, and add a NAT divert rule to the IP Firewall to direct network traffic to the appropriate computer. This also protects the internal network from unsolicited external connections.
- Enable DNS on the server, configured to cache lookups, to improve DNS response for the internal clients.

Before making of any of these settings, you have a chance to review the proposed changes before committing to them. It will overwrite any existing settings with those configured in the assistant.

You can make any additional changes to the service configuration using Server Admin. For any of the network services, see its appropriate section in this book for further instructions.

If you run the Gateway Setup Assistant again, it will overwrite any manual settings you may have made.

Using the Gateway Setup Assistant

The Gateway Setup Assistant can be accessed from two different locations. You can:

- Open `/Applications/Server/Gateway Setup`

or

- Choose `View > Gateway Setup` in Server Admin

Follow the directions in the assistant, and click Continue after each page. Read the final output carefully, and make sure you approve of the configuration before finalizing the settings.

Warning: Although Gateway Setup Assistant can be used to configure remote servers, you may accidentally cut off your administrative access to the remote server.

Sample Configurations

The following section contains a few sample configurations using the Gateway Setup Assistant. All of the configurations assume use the following information as an example:

- You have a static IP address allocated by the ISP (Internet Service Provider) for use by the server.
- The server is an XServe G5 with 2 built-in Ethernet ports as its network interfaces, Ethernet 1 (en0) and Ethernet 2 (en1), unless otherwise noted.

- The IP addresses to be used on your internal LAN are standard internal LAN IP addresses: 192.168.x.x

Connecting a Wired LAN to the Internet

You can use Gateway Setup Assistant to connect a wired LAN to the Internet. Your LAN may be made of any number of computers all connected to each other through Ethernet hubs and switches, but it has one point of contact with the Internet: the gateway.

After this process, all the computers on the LAN:

- Can get IP addresses and network settings configured via DHCP.
- Can access the Internet (as long as the gateway's connection to the Internet is present).
- Are not accessible to unauthorized network connections originating from the Internet.
- May be accessible via the Internet to authorized VPN clients (if configured).
- Benefit from DNS lookup caching in the gateway, speeding up DNS resolution.

To connect a wired LAN to the Internet:

- 1 Plug the connection to the Internet into the XServe's Built-In Ethernet 1 (en0) port.
- 2 Plug the connection to your LAN into the XServe's Built-In Ethernet 2 (en1) port.
- 3 Open the Gateway Setup Assistant.

You can open it from the /Applications/Server/ folder or via Server Admin's View menu.

Enter the address, administrator name, and password of the server you want to configure.

- 4 Designate Built-In Ethernet 1 as your WAN (Internet) interface.
- 5 Designate Built-In Ethernet 2 as your LAN (sharing) interface.

Your LAN interface is the one connected to your local network. All computers on the LAN will share the server's Internet connection through the server's WAN interface.

If your server has more than 1 interface available at this point (Ethernet 2 or Ethernet 3, and so on), choose those you wish to enable.

- 6 Choose whether to make this gateway a VPN entry point to your LAN.

If you choose to enable VPN, you'll need to have a "shared secret." A shared secret is a passphrase that all users must provide in order to make a secure connection to the VPN gateway. It should be a very secure passphrase, and not a password of any user or administrator on the gateway server.

For more information on VPN, see Chapter 6, "VPN Service," on page 99.

- 7 Inspect and confirm the changes.

Options:

You can fine-tune various settings from this base configuration. All further configuration is done using Server Admin.

For example, you can use Server Admin to set always give certain IP addresses to specific computers. You need to add static address mappings in the DHCP section's Settings tab. See Chapter 2, "DHCP Service," for more information.

Additionally, you can also change the IP Firewall settings to allow certain connections from the Internet to the LAN. You need to change the Firewall settings, opening up IP ports as desired, and configure port-forwarding (by editing UNIX files from the command line) to designate which computer on the LAN will accept the incoming traffic.

Connecting a Wired LAN And Wireless Clients To The Internet

You can use Gateway Setup Assistant to connect a wired LAN and wireless clients to the Internet. Your LAN may be made of any number of computers connected to each other through Ethernet hubs and switches, but it has one point of contact with the Internet: the gateway.

Your LAN must also have An AirPort Base Station to connect the wireless computers to the rest of the wired network. All your wireless clients must be able to connect to the AirPort base station's wireless network to be linked to the wired LAN.

After this process, the computers on the LAN and those connected to the AirPort Base Station:

- Can get IP addresses and network settings configured via DHCP.
- Can access the Internet (as long as the gateway's connection to the Internet is present).
- Are not accessible to unauthorized network connections originating from the wired connection to the Internet.
- May be accessible via the Internet to authorized VPN clients (if configured).
- Benefit from DNS lookup caching in the gateway, which speeds up DNS resolution.

To connect a wired LAN and wireless clients to the Internet:

- 1 Plug the connection to the Internet into the XServe's Built-In Ethernet 1 (en0) port.
- 2 Plug the connection to your LAN into the XServe's Built-In Ethernet 2 (en1) port.
- 3 Connect the AirPort Base Station's port (the WAN port, if there are 2) to the wired network.
- 4 Using the AirPort Admin Utility (or AirPort Setup Assistant) configure the base station to connect using Ethernet, and to get its own address using DHCP.
- 5 In the Network panel, make sure to uncheck "Distribute IP Addresses."
- 6 Click Update to change the base station settings.

7 Open the Gateway Setup Assistant.

You can open it from the /Applications/Server/ folder or via Server Admin's View menu.

Enter the address, administrator name, and password of the server you want to configure.

8 Designate Built-In Ethernet 1 as your WAN (Internet) interface.

9 Designate Built-In Ethernet 2 as your LAN (sharing) interface.

Your LAN interface is the one connected to your local network. All computers on the LAN will share the server's Internet connection through the server's WAN interface.

If your server has more than one interface available at this point (Ethernet 2 or Ethernet 3, and so on), choose those which you wish to enable.

10 Choose whether to make this gateway a VPN entry point to your LAN.

If you choose to enable VPN, you'll need to have a "shared secret." A shared secret is a passphrase that all users must provide in order to make a secure connection to the VPN gateway. It should be a very secure passphrase, and not a password of any user or administrator on the gateway server.

For more information on VPN, see Chapter 6, "VPN Service," on page 99.

11 Inspect and confirm the changes.

Options:

You can fine-tune various settings from this base configuration. All further configuration is done using Server Admin.

For example, you can use Server Admin to set always give certain IP addresses to specific computers. You need to add static address mappings in the DHCP section's Settings tab. See Chapter 2, "DHCP Service," for more information.

Additionally, you can also change the IP Firewall settings to allow certain connections from the Internet to the LAN. You need to change the Firewall settings, opening up IP ports as desired, and configure port-forwarding in the NAT panel to designate which computer on the LAN will accept the incoming traffic.

Connecting a Wireless LAN To The Internet

Connecting your wireless clients to the Internet through a Mac OS X Server gateway has some advantages over using the base stations built-in functions. The gateway can provide advanced IP Firewall control, DHCP allocation of static IP addresses, DNS caching, and incoming VPN connections to the LAN.

If you do not want or need these advanced functions, you can use the AirPort Base Station by itself to connect your wireless clients to the Internet, without a Mac OS X Server between the base station and the Internet.

To take advantage of the gateway's features, you use the base station as a bridge between your wireless clients and the gateway. Each client connects to the base station, and the base station sends the network traffic through the gateway. All your wireless clients must be able to connect to the AirPort Base Station's wireless network to be linked to the gateway.

After this process, the computers connected to the AirPort Base Station:

- Can get IP addresses and network settings configured via DHCP.
- Can access the Internet (as long as the gateway's connection to the Internet is present).
- Are not accessible to unauthorized network connections originating from the wired connection to the Internet.
- May be accessible via the Internet to authorized VPN clients (if configured).
- Benefit from DNS lookup caching in the gateway, speeding up DNS resolution.

To connect a wired LAN and wireless clients to the Internet:

- 1 Plug the connection to the Internet into the XServe's Built-In Ethernet 1 (en0) port.
- 2 Connect the AirPort Base Station's port (the WAN port, if there are 2) to the XServe's Built-In Ethernet 2 (en1) port.
- 3 Using the AirPort Admin Utility (or AirPort Setup Assistant) configure the base station to connect using Ethernet, and to get its own address using DHCP.
- 4 In the Network panel, make sure to uncheck "Distribute IP Addresses."
- 5 Click Update to change the base station settings.
- 6 Open the Gateway Setup Assistant.
You can open it from the /Applications/Server/ folder or via Server Admin's View menu. Enter the address, administrator name, and password of the server you want to configure.
- 7 Designate Built-In Ethernet 1 as your WAN (Internet) interface.
- 8 Designate Built-In Ethernet 2 as your LAN (sharing) interface.
Your LAN interface is the one connected to your local network. All computers on the LAN will share the server's Internet connection through the server's WAN interface. If your server has more than one interface available at this point (Ethernet 2 or Ethernet 3, etc.), choose those which you wish to enable.
- 9 Choose whether to make this gateway a VPN entry point to your LAN.
If you choose to enable VPN, you'll need to have a "shared secret." A shared secret is a passphrase that all users must provide in order to make a secure connection to the VPN gateway. It should be a very secure passphrase, and not a password of any user or administrator on the gateway server.

For more information on VPN, see Chapter 6, "VPN Service," on page 99.

- 10 Inspect and confirm the changes.

Options:

You can fine-tune various settings from this base configuration. All further configuration is done using Server Admin.

For example, you can use Server Admin to set always give certain IP addresses to specific computers. You need to add static address mappings in the DHCP section's Settings tab. See Chapter 2, "DHCP Service," for more information.

Additionally, you can also change the IP Firewall settings to allow certain connections from the Internet to the LAN. You need to change the Firewall settings, opening up IP ports as desired, and configure port-forwarding in the NAT panel to designate which computer on the LAN will accept the incoming traffic.

Dynamic Host Configuration Protocol (DHCP) service lets you administer and distribute IP addresses to client computers from your server. When you configure the DHCP server, you assign a block of IP addresses that can be made available to clients. Each time a client computer configured to use DHCP starts up, it looks for a DHCP server on your network. If a DHCP server is found, the client computer then requests an IP address. The DHCP server checks for an available IP address and sends it to the client computer along with a “lease period” (the length of time the client computer can use the address) and configuration information.

You can use the DHCP module in Server Admin to:

- Configure and administer DHCP service.
- Create and administer subnets.
- Configure DNS, LDAP, and WINS options for client computers.
- View DHCP address leases.

If your organization has more clients than IP addresses, you’ll benefit from using DHCP service. IP addresses are assigned on an as-needed basis, and when they’re not needed, they’re available for use by other clients. You can use a combination of static and dynamic IP addresses for your network if you need to. Read the next section for more information about static and dynamic allocation of IP addresses.

Organizations may benefit from the features of DHCP service, such as the ability to set Domain Name System (DNS) and Lightweight Directory Access Protocol (LDAP) options for client computers without additional client configuration.

Before You Set Up DHCP Service

Before you set up DHCP service, read this section for information about creating subnets, assigning static and dynamic IP addresses, locating your server on the network, and avoiding reserved IP addresses.

Creating Subnets

Subnets are groupings of computers on the same network that simplify administration. You can organize subnets any way that is useful to you. For example, you can create subnets for different groups within your organization or for different floors of a building. Once you have grouped client computers into subnets, you can configure options for all the computers in a subnet at one time instead of setting options for individual client computers. Each subnet needs a way to connect to the other subnets. A hardware device called a *router* typically connects subnets.

Assigning IP Addresses Dynamically

With dynamic allocation, an IP address is assigned for a limited period of time (the *lease time*) or until the client computer doesn't need the IP address, whichever comes first. By using short leases, DHCP can reassign IP addresses on networks that have more computers than available IP addresses. Leases are automatically renewed if the address isn't needed by any other computer.

Addresses allocated to Virtual Private Network (VPN) clients are distributed much like DHCP addresses, but they don't come out of the same range of addresses as DHCP. If you plan on using VPN, be sure to leave some addresses unallocated by DHCP for use by VPN. To learn more about VPN, see Chapter 6, "VPN Service," on page 99.

Using Static IP Addresses

Static IP addresses are assigned to a computer or device once and then don't change. You may want to assign static IP addresses to computers that must have a continuous Internet presence, such as web servers. Other devices that must be continuously available to network users, such as printers, may also benefit from static IP addresses.

Static IP addresses can be set up either by manually by entering the IP address on the computer or device that is assigned the address or by configuring DHCP to provide the same address to a specific computer or device on each request. DHCP-assigned addresses allow address configuration changes at the DHCP server rather than at each client. Manually configured static IP addresses avoid possible issues certain services may have with DHCP-assigned addresses and avoid the delay required for DHCP to assign an address.

Don't include manually assigned Static IP address ranges in the range distributed by DHCP.

It is possible to set up DHCP to always serve the same address to the same computer, giving you the benefits of static addresses and the benefits of centralized network configuration. See "Assigning Static IP Addresses Using DHCP" on page 30 for more information.

Locating the DHCP Server

When a client computer looks for a DHCP server, it broadcasts a message. If your DHCP server is on a different subnet from the client computer, you must make sure the routers that connect your subnets can forward the client broadcasts and the DHCP server responses. A relay agent or router on your network that can relay BootP communications will work for DHCP. If you don't have a means to relay BootP communications, you must place the DHCP server on the same subnet as your client.

Interacting With Other DHCP Servers

You may already have other DHCP servers on your network, such as AirPort Base Stations. Mac OS X Server can coexist with other DHCP servers as long as each DHCP server uses a unique pool of IP addresses. However, you may want your DHCP server to provide an LDAP server address for client autoconfiguration in managed environments. AirPort Base Stations can't provide an LDAP server address. Therefore, if you want to use the autoconfiguration feature, you must set up AirPort Base Stations in Ethernet-bridging mode and have Mac OS X Server provide DHCP service. If the AirPort Base Stations are on separate subnets, then your routers must be configured to forward client broadcasts and DHCP server responses as described previously. If you wish to provide DHCP service with AirPort Base Stations then you can't use the client autoconfiguration feature and you must manually enter LDAP server addresses at client workstations.

Using Multiple DHCP Servers on a Network

You can have multiple DHCP servers on the same network. However, it's important that they're configured properly to prevent interference with each other. Each server needs a unique pool of IP addresses to distribute.

Assigning Reserved IP Addresses

Certain IP addresses can't be assigned to individual hosts. These include addresses reserved for loopback and addresses reserved for broadcasting. Your ISP won't assign such addresses to you. If you try to configure DHCP to use such addresses, you'll be warned that the addresses are invalid, and you'll need to enter valid addresses.

Getting More Information on the DHCP Process

Mac OS X Server uses a daemon process called "bootpd" that is responsible for the DHCP Service's address allocation. You can learn more about bootpd and its advanced configuration options by accessing the bootpd man page by typing in to the Terminal:

```
man bootpd
```

Setting Up DHCP Service for the First Time

If you used the Setup Assistant to configure ports on your server when you installed Mac OS X Server, some DHCP information is already configured. You need to follow the steps in this section to finish configuring DHCP service. You can find more information about settings for each step in “Managing DHCP Service” on page 26.

Step 1: Create subnets

The following instructions show you how to create a pool of IP addresses that are shared by the client computers on your network. You create one range of shared addresses per subnet. These addresses are assigned by the DHCP server when a client issues a request.

See “Creating Subnets in DHCP Service” on page 27.

Step 2: Set up logs for DHCP service

You can log DHCP activity and errors to help you monitor requests and identify problems with your server.

DHCP service records diagnostic messages in the system log file. To keep this file from growing too large, you can suppress most messages by changing your log settings in the Logging pane of the DHCP service settings. For more information on setting up logs for DHCP service, see “Setting the Log Detail Level for DHCP Service” on page 32.

Step 3: Start DHCP service

See “Starting and Stopping DHCP Service” on page 26.

Managing DHCP Service

This section describes how to set up and manage DHCP service on Mac OS X Server. It includes starting service, creating subnets, and setting optional settings like LDAP or DNS for a subnet.

Starting and Stopping DHCP Service

Follow these steps when starting or stopping DHCP. You must have at least one subnet created and enabled.

To start or stop DHCP service:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Make sure at least one subnet and network interface is configured and selected.
- 3 Click Start Service or Stop Service.

When the service is turned on, the Stop Service button is available.

Creating Subnets in DHCP Service

Subnets are groupings of client computers on the same network that may be organized by location (different floors of a building, for example) or by usage (all eighth-grade students, for example). Each subnet has at least one range of IP addresses assigned to it.

To create a new subnet:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Click the Add (+) button.
- 5 Select the General tab.
- 6 Enter a descriptive name for the new subnet (optional).
- 7 Enter a starting and ending IP address for this subnet range.
Addresses must be contiguous, and they can't overlap with other subnets' ranges.
- 8 Enter the subnet mask for the network address range.
- 9 Choose the Network Interface from the pop-up menu.
- 10 Enter the IP address of the router for this subnet.
If the server you're configuring now is the router for the subnet, enter this server's internal LAN IP address as the router's address.
- 11 Define a lease time in hours, days, weeks, or months.
- 12 If you wish to set DNS, LDAP, or WINS information for this subnet, enter these now.
See "Setting the DNS Server for a DHCP Subnet" on page 29, "Setting LDAP Options for a Subnet" on page 29, and "Setting WINS Options for a Subnet" on page 30 for more information.
- 13 Click Save.

Changing Subnet Settings in DHCP Service

Use Server Admin to make changes to existing DHCP subnet settings. You can change IP address range, subnet mask, network interface, router, or lease time.

To change subnet settings:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Select a subnet.
- 5 Click the Edit (/) button.
- 6 Make the changes you want.

These changes can include adding DNS, LDAP, or WINS information. You can also redefine address ranges or redirect the network interface that responds to DHCP requests.

- 7 Click Save.

Deleting Subnets From DHCP Service

You can delete subnets and subnet IP address ranges when they will no longer be distributed to clients.

To delete subnets or address ranges:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select a subnet.
- 4 Click the Delete (-) button.
- 5 Click Save to confirm the deletion.

Disabling Subnets Temporarily

You can temporarily shut down a subnet without losing all its settings. This means no IP addresses from the subnet's range will be distributed on the selected interface to any client.

To disable a subnet:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Deselect "Enable" next to the subnet you want to disable.

Changing IP Address Lease Times for a Subnet

You can change how long IP addresses in a subnet are available to client computers.

To change the lease time for a subnet address range:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Select a subnet range and click the Edit (/) button.
- 5 Select the General tab.
- 6 Select a time scale from the Lease Time pop-up menu (hours, days, weeks, or months).
- 7 Enter a number in the Lease Time field.
- 8 Click Save.

Setting the DNS Server for a DHCP Subnet

You can decide which DNS servers and default domain name a subnet should use. DHCP service provides this information to the client computers in the subnet.

To set DNS options for a subnet:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Select a subnet and click the Edit (/) button.
- 5 Select the DNS tab.
- 6 Enter the default domain of the subnet.
- 7 Enter the primary and secondary name server IP addresses you want DHCP clients to use.
- 8 Click Save.

Setting LDAP Options for a Subnet

You can use DHCP to provide your clients with LDAP server information rather than manually configuring each client's LDAP information. The order in which the LDAP servers appear in the list determines their search order in the automatic Open Directory search policy.

If you have are using this Mac OS X Server as an LDAP master, the LDAP options will be prepopulated with the necessary configuration information. If your LDAP master server is another machine, you'll need to know the domain name or IP address of the LDAP database you want to use. You also will need to know the LDAP search base.

To set LDAP options for a subnet:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Select a subnet and click the Edit (/) button.
- 5 Click the LDAP tab.
- 6 Enter the domain name or IP address of the LDAP server for this subnet.
- 7 Enter the search base for LDAP searches.
- 8 Enter the LDAP port number, if you're using a non-standard port.
- 9 Select LDAP over SSL, if necessary.
- 10 Click Save.

Setting WINS Options for a Subnet

You can give additional information to client computers running Windows in a subnet by adding the Windows-specific settings to the DHCP supplied network configuration data. These Windows-specific settings allow Windows clients to browse their Network Neighborhood.

You must know the domain name or IP address of the WINS/NBNS primary and secondary servers (this is usually the IP address of the DHCP server itself), and the NBT node type (which is usually "broadcast"). The NBDD Server and the NetBIOS Scope ID are typically not used, but you may need to use them, depending on your Windows clients' configuration, and Windows network infrastructure.

To set WINS options for a subnet:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Select a subnet and click the Edit (/) button.
- 5 Click the WINS tab.
- 6 Enter the domain name or IP address of the WINS/NBNS primary and secondary servers for this subnet.
- 7 Enter the domain name or IP address of the NBDD server for this subnet.
- 8 Choose the NBT node type from the pop-up menu.
- 9 Enter the NetBIOS Scope ID.
- 10 Click Save.

Assigning Static IP Addresses Using DHCP

You can assign the same address to the same computers, if desired. This allows you to keep the ease of configuration of using DHCP, while allowing you to have some static servers or services.

In order to assign a the same IP address to the same computer, you will need the computer's Ethernet Address (sometimes called its MAC address, or hardware address). Each network interface has its own Ethernet Address.

Be aware that if you have computer that moves from being wired to the network to a wireless network, it uses two different Ethernet address, one for the wired connection, and one for the wireless connection.

To assign static IP addresses:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Click Static Maps.
- 4 Click the Add (+) button.
- 5 Enter the Ethernet Address of the computer which is to get a static address.
- 6 Enter the IP address you want to assign to it.
- 7 Enter the name of the computer.
- 8 Click OK.
- 9 Click Save.

Removing or Changing Static Address Maps

You can change the static mappings or remove them as needed.

To change the static address map:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Click Static Maps.
- 4 Select a mapping to edit or delete.
- 5 Click the Edit (/) button or the Delete (-) button.
- 6 If you are editing the mapping, make any changes you want, and click OK.
- 7 Click Save.

Monitoring DHCP Service

You'll need to monitor DHCP service. There are two main ways to monitor DHCP service. First, you can view the client list; second, you can monitor the log files generated by the service. You can use the service logs to help troubleshoot network problems. The following sections discuss these aspects of monitoring DHCP service.

Viewing the DHCP Status Overview

The status overview shows a simple summary of the DHCP service. It shows whether the service is running, how many clients it has, and when service was started. It also shows how many IP addresses are statically assigned from your subnets and the last time the client database was updated.

To see the overview:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click the Overview button.

Setting the Log Detail Level for DHCP Service

You can choose the level of detail you want to log for DHCP service.

- “Low (errors only)” will indicate conditions for which you need to take immediate action (for example, if the DHCP server can’t start up). This level corresponds to bootpd reporting in “quiet” mode, with the “-q” flag.
- “Medium (errors and warnings)” can alert you to conditions in which data is inconsistent, but the DHCP server is still able to operate. This level corresponds to default bootpd reporting.
- “High (all events)” will record all activity by the DHCP service, including routine functions. This level corresponds to bootpd reporting in “verbose” mode, with the “-v” flag.

To set up the log detail level:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.
- 4 Choose the logging option you want.
- 5 Click Save.

Viewing DHCP Log Entries

If you’ve enabled logging for DHCP service, you can check the system log for DHCP errors.

The log view is the system.log file filtered for “bootpd.” You can further filter the rules with the text filter box.

To see DHCP log entries:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Log.

Viewing the DHCP Client List

The DHCP Clients window gives the following information for each client:

- The IP address served to the client.
- The number of days of lease time left, until the time is less than 24 hours; then the number of hours and minutes.
- The DHCP client ID. This is usually, but not always, the same as the hardware address.
- The computer name.
- The Ethernet ID.

To view the DHCP client list:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Clients.

Click any column heading to sort the list by different criteria.

Common Network Configurations That Use DHCP

The following section contains some sample DHCP configurations for different network uses.

When you set up a private network, you choose IP addresses from the blocks of IP addresses that the IANA (Internet Assigned Numbers Authority) has reserved for private intranets:

- 10.0.0.0–10.255.255.255 (10/8 prefix)
- 172.16.0.0–172.31.255.255 (172.16/12 prefix)
- 192.168.0.0–192.168.255.255 (192.168/16 prefix)

Using DHCP to Provide IP Address Behind a NAT Gateway

You use DHCP to provide IP addresses to computers behind a NAT gateway. While not strictly necessary (NAT can be used with static IP addresses instead of DHCP), this allows easy configuration of client computers.

See “Linking a LAN to the Internet Through One IP Address” on page 93 for more information.

Workgroup Configuration

Let’s imagine you have a small workgroup with it’s own DHCP address group. You may have an IP-connected printer, a file server, and an Open Directory server (either on or off the subnet) for user management purposes. In order to use DHCP in this setting, you will need to already have:

- A working, configured firewall which allows LDAP and printer (IP Printing) connections.
See Chapter 4, “IP Firewall Service,” for more information.
- A working, configured Open Directory or LDAP server with users defined.
See the guide “*Mac OS X Server Open Directory Administration for Version 10.4 or Later*” and “*Mac OS X Server User Management for Version 10.4 or Later*” for more information.

Configuring DHCP, in this example, involves the use of static IP address mapping and additional client network settings. You could configure it like this:

- For a printer that must be given a static IP address, make sure the allocated DHCP address range does not include the truly static IP of the printer. If the printer can be configured to accept an address via DHCP, don’t worry about an overlap.
See “Using Static IP Addresses” on page 24 for more information.
- For a file server that needs to be allocated the same address all the time, use Mac OS X Server’s static IP mapping to always assign the same IP address to its Ethernet address.
See “Assigning Static IP Addresses Using DHCP” on page 30 for more information.

- For DHCP configuration, choose to set the LDAP options for DHCP clients. This automatically gives the clients their needed directory information. See “Setting LDAP Options for a Subnet” on page 29 for more information.
- For client configuration on Mac OS X Clients, make sure the IPv4 configuration method in the Network pane of System Preferences is set to “DHCP.”

This configuration will allow the computers on the network to be managed via an LDAP or Open Directory server, getting all their networking configuration from DHCP. They can have access to truly static IP address or consistently assigned IP addresses on the same network. Additionally, you get centralized configuration for all the computer clients.

Student Lab Configuration

The student lab configuration is very much like the workgroup configuration, but it adds an additional service that uses DHCP: Netboot. In addition to DHCP for centralized networking configuration, Netboot standardizes start-up environments by having each client computer start up from a disk image on a central Netboot server.

The configuration would be like the “Workgroup Configuration” on page 33, with the following exceptions:

- There may or may not be static-address resources.
This depends on the lab composition, of course. You might have a class printer or file server, but if you use a mobile cart that moves from classroom to classroom, you won’t tote a server and printer along to each class.
- NetBoot is enabled and configured, along with the firewall settings to support it. Any client on the network can be set to start up from the NetBoot server. New computers can be deployed by setting the startup disk of the computer to the NetBoot image. No additional configuration necessary, and computers can be re-purposed easily, since the hard drive can remain unchangeable.

This configuration will allow the computers on the network to be managed via an LDAP or Open Directory server, getting all their networking configuration from DHCP. The computing environment is also centrally configured for all the computer clients. New clients can be added or swapped out with minimal effort.

Coffee Shop Configuration

The “coffee shop configuration” isn’t specifically about a coffee shop. It is a type of configuration for a purely dynamic addressing environment, with no user management, and no services provided other than web access, DNS access, and possibly some other service. Its characterized by lots of mobile users who pass through, use the Internet access, and move on. This configuration could easily be used in real-life situations like a college-commons wireless network, or a wired courtesy office for visiting consultants.

Warning: Make sure any sensitive information you may have on your LAN is well protected behind an additional Firewall on another network, if you host temporary, unauthenticated users.

In order to use DHCP in this setting, you will need to already have:

- A working, configured firewall which allows web access outbound traffic, and DNS outbound lookups only. You may want to place this network outside your firewall, and make sure the DHCP allocated IP addresses’ network traffic is strictly controlled and monitored.

See Chapter 4, “IP Firewall Service,” for more information.

You might want to configure the DHCP service like this:

- *Make networking configuration automatic.* Set the DHCP clients to get all possible network configuration via DHCP.
- *Don’t set options they shouldn’t have.* Don’t give the DHCP clients additional information about your organization via LDAP information. You may want to allow Windows clients to have additional network options.
See “Setting WINS Options for a Subnet” on page 30 for more information.
- *Limit resource use.* Many users can lead to a lot of bandwidth use, so you could reduce the number of DHCP clients that can be connected simultaneously by allowing only a smaller number of addresses to be allocated.
See “Creating Subnets in DHCP Service” on page 27 for more information.
- *Keep address turnover high.* You’ll want to make the lease times on addresses as short as possible. This way, as users come and go from the network, the addresses can be reallocated as quickly as possible.
See “Creating Subnets in DHCP Service” on page 27 for more information.
- Monitor your traffic. You may want to keep a closer eye on DHCP connections and clients, firewall rule packet logging, or other monitoring tools. Open access points can be a liability if not guarded vigilantly.

Where to Find More Information

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at:

www.ietf.org/rfc.html

For details about DHCP, see RFC 2131.

For more information on `bootpd` and its advanced configuration options, see `bootpd`'s man page in the Terminal:

```
man bootpd
```

When your clients want to connect to a network resource such as a web or file server, they typically request it by its domain name (such as `www.example.com`) rather than by its IP address (such as `192.168.12.12`). The Domain Name System (DNS) is a distributed database that maps IP addresses to domain names so your clients can find the resources by name rather than by numerical address.

A DNS server keeps a list of domain names and the IP addresses associated with each name. When a computer needs to find the IP address for a name, it sends a message to the DNS server (also known as a *name server*). The name server looks up the IP address and sends it back to the computer. If the name server doesn't have the IP address locally, it sends messages to other name servers on the Internet until the IP address is found.

Setting up and maintaining a DNS server is a complex process. Therefore many administrators rely on their Internet Service Provider (ISP) for DNS services. In this case, you only have to configure your network preferences with the name server IP address provided by your ISP.

If you don't have an ISP to handle DNS requests for your network and any of the following is true, you need to set up DNS service:

- You don't have the option to use DNS from your ISP or other source.
- You plan on making frequent changes to the namespace and want to maintain it yourself.
- You have a mail server on your network and you have difficulties coordinating with the ISP that maintains your domain.
- You have security concerns regarding giving your network's computer names and addresses to an outside organization (your ISP).

Mac OS X Server uses Berkeley Internet Name Domain (BIND v.9.2.2) for its implementation of DNS protocols. BIND is an open-source implementation and is used by the majority of name servers on the Internet.

Before You Set Up DNS Service

This section contains information you should consider before setting up DNS on your network. The issues involved with DNS administration are complex and numerous. You should not set up DNS service on your network unless you're an experienced DNS administrator.

You should consider creating a mail alias called "hostmaster" that receives mail and delivers it to the person that runs the DNS server at your site. This allows users and other DNS administrators to contact you regarding DNS problems.

DNS and BIND

You should have a thorough understanding of DNS before you attempt to set up your own DNS server. A good source of information about DNS is *DNS and BIND, 4th edition*, by Paul Albitz and Cricket Liu (O'Reilly and Associates, 2001).

Note: Apple can help you locate a network consultant to implement your DNS service. You can contact Apple Professional Services and Apple Consultants Network on the web at www.apple.com/services/ or www.apple.com/consultants.

Setting Up Multiple Name Servers

You should set up at least one primary and one secondary name server. That way, if the primary name server unexpectedly shuts down, the secondary name server can continue to provide service to your users. A secondary server gets its information from the primary server by periodically copying all the domain information from the primary server.

Once a name server learns a name/address pair of a host in another domain (outside the domain it serves), the information is cached, which ensures that IP addresses for recently resolved names are stored for later use. DNS information is usually cached on your name server for a set time, referred to as a *time-to-live* (TTL) value. When the TTL for a domain name/IP address pair has expired, the entry is deleted from the name server's cache and your server will request the information again as needed.

Setting Up DNS Service for the First Time

If you're using an external DNS name server and you entered its IP address in the Setup Assistant, you don't need to do anything else. If you're setting up your own DNS server, follow the steps in this section.

Step 1: Register your domain name

Domain name registration is managed by a central organization, the Internet Assigned Numbers Authority (IANA). IANA registration makes sure domain names are unique across the Internet. (See www.iana.org for more information.) If you don't register your domain name, your network won't be able to communicate over the Internet.

Once you register a domain name, you can create subdomains within it as long as you set up a DNS server on your network to keep track of the subdomain names and IP addresses.

For example, if you register the domain name “example.com,” you could create subdomains such as “host1.example.com,” “mail.example.com,” or “www.example.com.” A server in a subdomain could be named “primary.www.example.com,” or “backup.www.example.com.” The DNS server for example.com keeps track of information for its subdomains, such as host (or computer) names, static IP addresses, aliases, and mail exchangers. If your ISP handles your DNS service, you’ll need to inform them of any changes you make to your namespace, including adding subdomains.

The range of IP addresses for use with a given domain must be clearly defined before setup. These addresses are used exclusively for one specific domain (never by another domain or subdomain). The range of addresses should be coordinated with your network administrator or ISP.

Step 2: Learn and plan

If you’re new to working with DNS, learn and understand DNS concepts, tools, and features of Mac OS X Server and BIND. See “Where to Find More Information” on page 58.

Then plan your Domain Name System Service. You may consider the following questions when planning:

- Do you even need a local DNS server? Does your ISP provide DNS service? Could you use Multicast DNS names instead?
- How many servers will you need for the anticipated load? How many servers will you need for backup purposes? For example, you should designate a second or even third computer for backup DNS service.
- What is your security strategy to deal with unauthorized use?
- How often should you schedule periodic inspections or tests of the DNS records to verify data integrity?
- How many services or devices (like an intranet website or a network printer) are there that will need a name?

There are two ways to configure DNS service on Mac OS X Server. First, and recommended, you can use Server Admin to set up DNS service. For more information, see “Managing DNS Service” on page 41 for instructions.

The second way to configure DNS is by editing the BIND configuration file. BIND is the set of programs used by Mac OS X Server that implements DNS. One of those programs is the *name daemon*, or *named*. To set up and configure BIND, you need to modify the configuration file and the zone file.

The configuration file is located in this file:

```
/etc/named.conf
```

The zone file name is based on the name of the zone. For example, the zone file "example.com" is located in this file:

```
/var/named/example.com.zone
```

If you edit named.conf to configure BIND, make sure that you don't change the controls statement `inet` settings. Otherwise, Server Admin will be unable to retrieve status information for DNS.

The `inet` settings should look like this:

```
controls {
    inet 127.0.0.1 port 54 allow {any;}
    keys { "rndc-key"; };
};
```

Step 3: Configure basic DNS settings

See "Managing DNS Service" on page 41 for more information.

You should decide if you want to allow Recursion or Zone Transfers.

Step 4: Create a DNS Zone

Use Server Admin to set up DNS zones. See "Managing DNS Zones" on page 42 for instructions. After adding a primary zone, Server Admin automatically creates an NS record with the same name as the Source of Authority (SOA). For each zone that you create, Mac OS X Server creates a reverse lookup zone. Reverse lookup zones translate IP addresses to domain names, rather than normal lookups which translate domain names to IP addresses.

Step 5: Add DNS machine records to the zone

Use Server Admin to add records to your Zone. Create an Address record for every computer or device (printer, file server, and so on) that has a static IP address and needs a name. Various DNS zone records are created from the DNS machine entries. See "Managing DNS Machine Records" on page 46 for instructions.

Step 6: Set up a mail exchange (MX) record (optional)

If you provide mail service over the Internet, you need to set up an MX record for your server. See "Setting Up MX Records" on page 53 for more information.

Step 7: Configure IP Firewall

You will need to configure your firewall to make sure your DNS service is protected from attack, and accessible to your clients.

See Chapter 4, "IP Firewall Service," for more information on configuring IP Firewall.

Step 8: Start DNS service

Mac OS X Server includes a simple interface for starting and stopping DNS service.

See “Starting and Stopping DNS Service” on page 41 for more information.

Managing DNS Service

Mac OS X Server provides a simple interface for starting and stopping DNS service as well as viewing logs and status. Basic DNS settings can be configured with Server Admin. More advanced features require configuring BIND from the command-line, and are not covered here.

Starting and Stopping DNS Service

Use this procedure to start or stop DNS service. Remember to restart the DNS service whenever you make changes to the DNS service in Server Admin.

To start or stop DNS service:

- 1 In Server Admin, choose DNS from the Computers & Services list.
- 2 Make sure you have at least one Zone and its reverse lookup zone created and fully configured.
- 3 Click Start Service or Stop Service.

The service may take a moment to start (or stop).

Enabling or Disabling Zone Transfers

In the Domain Name System, zone data is replicated among authoritative DNS servers by means of the “zone transfer.” Secondary DNS servers (“secondaries”) use zone transfers to acquire their data from primary DNS servers (“primaries”). Zone transfers must be enabled to use secondary DNS servers.

To enable or disable zone transfer:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select or deselect Allow Zone Transfers as needed.

Enabling or Disabling Recursion

Recursion is a process to fully resolve domain names into IP addresses. Users' applications depend on the DNS server to perform this function. Other DNS servers that query yours don't have to perform the recursion.

To prevent malicious users from altering the primary zone's records ("cache poisoning"), or allowing unauthorized use of the server for DNS service, you can disable recursion. However, if you stop it, your own users won't be able to use your DNS service to look up any names outside of your zones.

You should disable recursion only if no clients are using this DNS server for name resolution and no servers are using it for forwarding.

To enable or disable recursion:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select or deselect Recursion as needed.

Selecting Recursion allows it. Deselecting Recursion disallows it.

If you choose to enable recursion, consider disabling it for external IP addresses, but enabling it for LAN IP addresses, by editing BIND's `named.conf` file. See BIND's documentation for more information.

Managing DNS Zones

Zones are the basic organizational unit of the Domain Name System. Zones contain records and are defined by how they acquire those records, and how they respond to DNS requests. There are normally three basic kinds of zones (and a few not covered here):

Primary

A primary zone has the master copy of the zone's records, and provides authoritative answers to lookup requests.

Secondary

A secondary zone is a copy of a primary zone stored on a secondary name server. Each secondary zone keeps a list of primary servers that it contacts to receive updates to records in the primary zone. Secondaries must be configured to request the copy of the primary zone's data. Secondary zones use zone transfers to get copies of the primary zone data. Secondary name servers can take lookup requests like primary servers. By using several secondary zones linked to one primary, you can distribute DNS query loads across several computers and ensure lookup requests are answered when the primary name server is down.

Secondary zones also have a refresh interval also. It determines how often secondary zones check for changes from the primary zone. You can change the zone refresh interval by using BIND's configuration file. See BIND's documentation for more information.

Forward

A forward zone directs all lookup requests for that zone to other DNS servers. Forward zones don't do zone transfers. Often, forward zone servers are used to provide DNS services to a private network behind a firewall. In this case, the DNS server must have access to the Internet and a DNS server outside the firewall. Finally, forward zones cache responses to the queries that they pass on. This can improve the performance of lookups by clients which use the forward zone.

Server Admin does not support creation or modification of a forward zone. In order to create a forward zone, you will need to configure BIND manually on the command line. See BIND's documentation for further detail.

Adding a Primary Zone

A primary zone has the master copy of the zone's records and provides authoritative answers to lookup requests. After adding a primary zone, Server Admin automatically creates an NS record with the same name as the Source of Authority (SOA).

To add a primary zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Click the Add (+) button beneath the Zones list.
- 5 Enter a zone name.

The zone name is domain name.

- 6 Enter the hostname of the domain's SOA.

If this computer will be the authoritative name server for the domain, enter the computer's hostname. For example, "ns.example.com."

- 7 Enter the zone server's IP address.
- 8 Enter the email address of the zone's administrator.
- 9 Enter the amount of time the zone is valid.

This is the zone's Time to Live (TTL). It determines how long any query responses information can remain cached in remote DNS systems before re-querying the authoritative server.

- 10 Click Save.

Adding a Secondary Zone

A secondary zone is a copy of a primary zone stored on a secondary name server. Each secondary zone keeps a list of primary servers that it contacts to receive updates to records in the primary zone. Secondaries must be configured to request the copy of the primary zone's data. Secondary zones use zone transfers to get copies of the primary zone data. Secondary name servers can take lookup requests like primary servers.

To add a secondary zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Secondary Zones tab.
- 4 Click the Add (+) button beneath the Zones list.
- 5 Enter a zone name.
This is the fully qualified domain name of the secondary server.
- 6 Click the Add (+) button.
- 7 Enter the IP addresses for the primary servers for this secondary zone.
- 8 Click OK
- 9 Click Save.

Duplicating a Zone

You can create a copy of an existing zone on the same computer. You could use this to speed up configuration of multiple zones, or multiple domain names for a single physical LAN.

To duplicate a zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Click the Duplicate button beneath the Zones list.
- 5 If desired, double-click the newly duplicated zone to change the zone information.
- 6 Click Save.

Modifying a Zone

This section describes modifying a zone's type and settings but not modifying the records within a zone. You may need to change a zone's administrator address, type, or domain name.

To modify a zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Click the Edit (/) button beneath the Zones list.
- 5 Change the zone name, type, or administrator email address as needed.

For more information on zone types, see "Managing DNS Zones" on page 42.

- 6 Click OK, and click Save.

Deleting a Zone

The section describes how to delete an existing zone. This deletes the zone and all the records associated with it.

To delete a zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the zone which is to be deleted.
- 5 Click the Delete (-) button beneath the Zones list.
- 6 Click Save to confirm the deletion.

Using an Existing Zone File

You may already have a BIND zone file from a DNS server of another platform. Instead of entering all the information in to Server Admin by hand, you can use the zone file directly with Mac OS X Server.

Using an existing zone file requires root access permissions to the BIND configuration file (`/etc/named.conf`), the working zone directory (`/var/named/`), a basic knowledge of BIND 9, and a facility with the Terminal application. Otherwise, it is strongly advised you use Server Admin's DNS tools.

To import a zone file:

- 1 Add the zone directive to BIND's configuration file, /etc/named.conf

You'll need root privileges to edit named.conf.

For a zone "xyz.com" described in a zone file "db.xyz.com" in working zone directory "/var/named/", the directive might look something like this:

```
zone "xyz.com" IN {           // Forward lookup zone for xyz.com
    type master;             // It's a primary zone
    file "db.xyz.com";      // Zone info stored in /var/named/db.xyz.com
    allow-update { none; };
};
```

- 2 Make sure the zone file is added to the working zone directory (/var/named/).
- 3 Restart the DNS service using Server Admin.

Managing DNS Machine Records

Each zone contains a number of records. These records are requested when a client computer needs to translate a domain name (like www.example.com) to an IP number. Web browsers, email clients, and other network applications rely on a zone's records to contact the appropriate server. The primary zone's records will be queried by others across the Internet so they can connect to your network services. There are several kinds of DNS records. The records which are available for configuration by Server Admin's user interface are:

- *Address (A)*: Stores the IP address associated with a domain name.
- *Canonical Name (CNAME)*: Stores an alias in connection with the "real name" of a server. For example, mail.apple.com might be an alias for a computer with a "real" canonical name of MailSrv473.apple.com.
- *Mail Exchanger (MX)*: Stores the domain name of the computer that is used for email in a zone.
- *Name Server (NS)*: Stores the authoritative name server for a given zone.
- *Pointer (PTR)*: Stores the domain name of a given IP address (reverse lookup).
- *Text (TXT)*: Stores a text string as a response to a DNS query.
- *Service (SRV)*: Stores information about what services a computer provides.
- *Hardware Info (HINFO)*: Stores information about a computer's hardware and software.

If you need access to other kinds of records, you'll need to edit BIND's configuration files manually. Please see BIND's documentation for details.

Mac OS X Server simplifies the creation of all these records by focusing on the computer being added to the zone rather than the records themselves. As you add a computer record to a zone, Mac OS X Server creates all the appropriate zone records that resolve to a certain computer address.

With this model, you can focus on what your computers *do* in your domain, rather than *which* record types apply to its functions.

Adding a Machine Record to a DNS Zone

You need to add records for each computer for which the DNS primary zone has responsibility. You should not add records for computers that this zone doesn't control. Machine records are tied to its IP address. Therefore, there can be only one machine per IP address because there can be no duplicate IP addresses within a zone.

To add a DNS machine record:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the zone to which this record will be added.
- 5 Click the Edit (/) button beneath the zone list.
- 6 Select the Machines tab.
- 7 Click the Add (+) button beneath the machines list.
- 8 Enter the computer's IP address.
- 9 Enter the computer's hostname.

This field is the basis for the computer's A record. Reverse lookup Pointer records are automatically created for the computer.

Beneath the hostname, you will see what will be the computer's Fully Qualified Domain Name.

- 10 Click the Add (+) button by the Alias box to add other names as you want for this computer.

This field is the basis for the computer's CNAME records. Reverse lookup Pointer records are automatically created for the computer.

Add as many aliases as you want.

- 11 If the computer is a mail server for the zone, check the indicated box.

This field is the basis for the computer's MX record.

If you check the box, set a mail server precedence number. Delivering mail servers try to deliver mail at lower numbered mail servers first. See "Setting Up MX Records" on page 53 for more information.

- 12 Enter any information about the computer's hardware and software in the appropriate boxes.

This field is the basis for the computer's HINFO record.

- 13 Enter any comments about the computer in the Comments box.

This field is the basis for the computer's TXT record.

You can store almost any 7-bit ASCII text string in the comments box (up to 255 ASCII characters). For example, you might include the physical location of the computer (for example, Upstairs server closet B) or the computer's owner (for example, John's Computer) or any other information you may want to keep about the computer.

- 14 Click OK, and click Save.

Modifying a Machine Record in a DNS Zone

If you make frequent changes to the namespace for the domain, you'll need to update the DNS records as often as that namespace changes. Upgrading hardware or adding to a domain name might require updating the DNS records as well.

You can choose to duplicate a record, and then edit it, saving some configuration time.

To modify a record:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the zone which has the computer record to be edited.
- 5 Click the Edit (/) button beneath the zone list.
- 6 Select the Machines tab.
- 7 Select the record to be edited.
- 8 Click the Edit (/) button beneath the machines list.
- 9 Modify the record as needed.
- 10 Click OK.

Deleting a Machine Record From a DNS Zone

You should delete records whenever a computer is no longer associated with a domain name or usable address.

To delete a record:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the zone from which this record will be deleted.
- 5 Click the Edit (/) button beneath the zone list.
- 6 Select the Machines tab.

- 7 Select the record to be deleted.
- 8 Click the Delete (-) button beneath the Records list.
- 9 Click Save to confirm the deletion.

Monitoring DNS

You may want to monitor DNS status to troubleshoot name resolution problems, check how often the DNS service is used, or even check for unauthorized or malicious DNS service use. This section discusses common monitoring tasks for DNS service.

Viewing DNS Service Status

You can check the DNS Status window to see:

- Whether the service is running.
- The version of BIND (the underlying software for DNS) that is running.
- When the service was started and stopped.
- The number of zones allocated.
- Whether logging is on or off.

To view DNS service status:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click the Overview button for general DNS service information.

Viewing DNS Log Entries

DNS service creates entries in the system log for error and alert messages. The log view is the named.log. You can further filter the rules with the text filter box.

To see DNS log entries:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Log.
- 3 Use the filter field to further focus your view of the log entries.

Changing DNS Log Detail Levels

You can change the detail level of the DNS service log. You may want a highly detailed log for debugging, or a less detailed log that only shows critical warnings.

To change the log detail level:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.

- 4 Choose the detail level from the Log Level pop-up menu.

The possible log levels are:

- Critical (less detailed)
- Error
- Warning
- Notice
- Information
- Debug (most detailed)

Changing DNS Log File Location

You can change the location of the DNS service log. You may want to put it somewhere other than the default path.

To change the log detail level:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.
- 4 Enter the desired path for the file path for the DNS service log, or select a path using the Browse button.

If no path is entered, the default location is /Library/Logs/named.log.

Securing the DNS Server

DNS servers are targeted by malicious computer users (commonly called “hackers”) in addition to other legitimate Internet servers. There are several kinds of attacks that DNS servers are susceptible to. By taking extra precautions, you can prevent the problems and downtime associated with malicious users. There are several kinds of security hacks associated with DNS service:

- DNS Spoofing.
- Server Mining.
- DNS Service Profiling.
- Denial of Service (DoS).
- Service Piggybacking.

DNS Spoofing

DNS spoofing is adding false data into the DNS Server’s cache. This allows hackers to do any of the following:

- Redirect real domain name queries to alternative IP Addresses.

For example, a falsified A record for a bank could point a computer user’s browser to a different IP address that is controlled by the hacker. A duplicate website could fool him or her into giving their bank account numbers and passwords to the hacker.

Also, a falsified mail record could allow a hacker to intercept mail sent to or from a domain. If the hacker also forwards those emails to the correct mail server after copying them, this can go undetected indefinitely.

- Prevent proper domain name resolution and access to the Internet.
This is the most benign of DNS spoof attacks. It merely makes a DNS server appear to be malfunctioning.

The most effective method to guard against these attacks is vigilance. This includes maintaining up-to-date software as well as auditing your DNS records regularly. As exploits are found in the current version of BIND, the exploit is patched and a Security Update is made available for Mac OS X Server. Apply all such security patches. Regular audits of your DNS records can help prevent these attacks.

Server Mining

Server mining is the practice of getting a copy of a complete primary zone by requesting a zone transfer. In this case, a hacker pretends to be a secondary zone to another primary zone and requests a copy of all of the primary zone's records.

With a copy of your primary zone, the hacker can see what kinds of services a domain offers, and the IP address of the servers that offer them. He or she can then try specific attacks based on those services. This is reconnaissance before another attack.

To defend against this attack, you need to specify which IP addresses are allowed to request zone transfers (your secondary zone servers) and disallow all others. Zone transfers are accomplished over TCP on port 53. The method of limiting zone transfers is blocking zone transfer requests from anyone but your secondary DNS servers.

To specify zone transfer IP addresses:

- Create a firewall filter that allows only IP addresses inside your firewall to access TCP port 53.

Follow the instructions in "Creating an Advanced IP Firewall Rule" in Chapter 4, "IP Firewall Service." Use the following settings:

- Allow packet.
- Port 53.
- TCP protocol.
- Source IP is the IP address of your secondary DNS server.
- Destination IP is the IP address of your primary DNS server.

DNS Service Profiling

Another common reconnaissance technique used by malicious users is to profile your DNS Service. First a hacker makes a BIND version request. The server will report what version of BIND is running. He or she then compares the response to known exploits and vulnerabilities for that version of BIND.

To defend against this attack, you can configure BIND to respond with something other than what it is.

To alter BIND's version response:

- 1 Launch a command-line text editor (like vi, emacs, or pico).
- 2 Open named.conf for editing.
- 3 Add the following to the "options" brackets of the configuration file.

```
version "[your text, maybe 'we're not telling!']";
```

- 4 Save the config file.

Denial of Service (DoS)

This kind of attack is very common and easy to do. A hacker sends so many service requests and queries that a server uses all of its processing power and network bandwidth trying to respond. The hacker prevents legitimate use of the service by overloading it.

It is difficult to prevent this type of attack before it begins. Constant monitoring of the DNS service and server load allows an administrator to catch the attack early and mitigate its damaging effect.

The easiest way to guard against this attack is to block the offending IP address with your firewall. See "Creating an Advanced IP Firewall Rule" on page 70. Unfortunately, this means the attack is already underway and the hacker's queries are being answered and the activity logged.

Service Piggybacking

This attack is not often done by malicious intruders, but common Internet users. They may feel that their DNS response time with their own Internet Service Provider is too slow. They learn this trick from other users. The Internet users will configure their computer to query another DNS server instead of their own ISP's DNS servers. Effectively, there will be more users accessing the DNS server than have been planned for.

You can guard against this by limiting or disabling DNS Recursion. If you plan to offer DNS service to your own LAN users, they need recursion to resolve domain names, but you don't want to provide this service to any Internet users.

To prevent recursion entirely, see "Enabling or Disabling Recursion" on page 42.

The most common balance is allowing recursion for requests coming from IP addresses within your own range, but denying recursion to external addresses. BIND allows you to specify this in its configuration file, `named.conf`. Edit your `named.conf` file to include the following:

```
options {  
    ...  
    allow-recursion{  
        127.0.0.0/8;  
        [your internal IP range of addresses, like 192.168.1.0/27];  
    };  
};
```

Please see BIND's documentation for further information.

Common Network Administration Tasks That Use DNS Service

The following sections illustrate some common network administration tasks that require DNS service.

Setting Up MX Records

If you plan to provide mail service on your network, you must set up DNS so that incoming mail is sent to the appropriate mail host on your network. When you set up mail service, you define a series of hosts, known as *mail exchangers* or *MX hosts*, with different priorities. The host with the highest priority gets the mail first. If that host is unavailable, the host with the next highest priority gets the mail, and so on.

For example, let's say the mail server's host name is "reliable" in the "example.com" domain. Without an MX record, the users' mail addresses would include the name of your mail server computer, like this:

```
user-name@reliable.example.com
```

If you want to change the mail server or redirect mail, you must notify potential senders of a new address for your users. Or, you can create an MX record for each domain that you want handled by your mail server and direct the mail to the correct computer.

When you set up an MX record, you should include a list of all possible computers that can receive mail for a domain. That way, if the server is busy or down, mail is sent to another computer. Each computer on the list is assigned a priority (“precedence”) number. The one with the lowest number is tried first. If that computer isn’t available, the computer with the next lowest number is tried, and so on. When a computer is available, it holds the mail and sends it to the main mail server when the main server becomes available, and then the server delivers the mail. A sample list might look like this:

example.com

10 reliable.example.com

20 our-backup.example.com

30 last-resort.example.com

MX records are used for outgoing mail, too. When your mail server sends mail, it looks at the MX records to see whether the destination is local or somewhere else on the Internet. Then the same process happens in reverse. If the main server at the destination is not available, your mail server tries every available computer on that destination’s MX record list, until it finds one that will accept the mail.

Note: If you don’t enter the MX information into your DNS server correctly, mail won’t work.

Configuring DNS for Mail Service

Configuring DNS for mail service is creating a Mail Exchange (MX) records in DNS for your mail servers. If you have an Internet Service Provider (ISP) that provides you with DNS service, you’ll need to contact the ISP so that they can enable your MX records. Only follow these steps if you provide your own DNS Service.

You may need to set up multiple servers for redundancy. If this is the case, you’ll need to create an MX record for each auxiliary server.

To enable MX records for your mail server:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the zone to which this record will be added.
- 5 Click the Edit (/) button beneath the zone list.
- 6 Select the Machines tab.
- 7 Click the Add (+) button beneath the machines list.
- 8 Enter the computer’s IP address.

- 9 Enter the computer's hostname.

This field is the basis for the computer's CNAME and first A record. Reverse lookup Pointer records are automatically created for the computer.

Beneath the hostname, you will see what will be the computer's Fully Qualified Domain Name.

- 10 Click the Add (+) button by the Alias box to add other names as you want for this computer.

This field is the basis for the computer's additional A records. Reverse lookup Pointer records are automatically created for the computer.

Add as many aliases as you want.

- 11 Check mailer server box, labeled "This Machine Is A Mail Server For The Zone."

This field is the basis for the computer's MX record.

- 12 Enter a mail server precedence number.

Delivering mail servers try to deliver mail at lower numbered mail servers first.

- 13 Enter any information about the computer's hardware and software in the appropriate boxes.

This field is the basis for the computer's HINFO record.

- 14 Enter any comments about the computer in the Comments box.

This field is the basis for the computer's TXT record.

You can store almost any text string in the comments box. For example, you might include the physical location of the computer (for example, Upstairs server closet B) or the computer's owner (for example, John's Computer) or any other information you may want to keep about the computer.

- 15 Click OK.

- 16 Repeat Steps 7 through 15 for each mail server, making sure each has a distinct precedence number.

- 17 Click Save.

Setting Up Namespace Behind a NAT Gateway

If you're behind a Network Address Translation (NAT) gateway, you have a special set of IP addresses that are only usable within the NAT environment. If you were to assign a domain name to these addresses outside of the NAT gateway, none of your domain names would resolve to the correct computer. See Chapter 5, "NAT Service," on page 87 for more information about NAT.

You can, however, run a DNS service behind the gateway, assigning host names to the NAT IP addresses. This way, if you're behind the NAT gateway, you can enter domain names rather than IP addresses to access servers, services, and workstations. Your DNS server should also have a Forwarding zone to send DNS requests outside of the NAT gateway to allow resolution of names outside the routed area. Your clients' networking settings should specify the DNS server behind the NAT gateway. The process of setting up one of these networks is the same as setting up a private network. See "Linking a LAN to the Internet Through One IP Address" on page 93 for more information.

If you choose to do this, names entered by users outside the NAT gateway won't resolve to the addresses behind it. You should set the DNS records outside the NAT-routed area to point to the NAT gateway, and use NAT port forwarding to access computers behind the NAT gateway. For more information on port forwarding, see "Configuring Port Forwarding" on page 90.

Mac OS X's Multicast DNS feature allows you to use hostnames on your local subnet that end with the ".local" suffix without having to enable DNS. Any service or device that supports Multicast DNS allows the use of user-defined namespace on your local subnet without setting up and configuring DNS.

Network Load Distribution (aka Round Robin)

BIND allows for simple load distribution using an address-shuffling method called *round robin*. You set up a pool of IP addresses for several hosts mirroring the same content, and BIND cycles the order of these addresses as it responds to queries. Round robin has no capability to monitor current server load or processing power. It simply cycles the order of an address list for a given host name.

You enable round robin by adding multiple IP address entries for a given hostname. For example, suppose you want to distribute web server traffic between three servers on your network that all mirror the same content. Suppose the servers have the IP addresses 192.168.12.12, 192.168.12.13, and 192.168.12.14. You would add three machine records with three IP Addresses, each with the same domain name.

When the DNS service encounters multiple entries for one host, its default behavior is to answer queries by sending out this list in a cycled order. The first request gets the addresses in the order A, B, C. The next request gets the order B, C, A, then C, A, B, and so on. You may want the zone's *time-to-live* (TTL) number to be fairly short to mitigate the effects of local caching.

Setting Up a Private TCP/IP Network

If you have a local area network that has a connection to the Internet, you must set up your server and client computers with IP addresses and other information that's unique to the Internet. You obtain IP addresses from your Internet service provider (ISP).

If it's unlikely that your local area network will ever be connected to the Internet and you want to use TCP/IP as the protocol for transmitting information on your network, it's possible to set up a "private" TCP/IP network. When you set up a private network, you choose IP addresses from the blocks of IP addresses that the IANA (Internet Assigned Numbers Authority) has reserved for private intranets:

- 10.0.0.0–10.255.255.255 (10/8 prefix)
- 172.16.0.0–172.31.255.255 (172.16/12 prefix)
- 192.168.0.0–192.168.255.255 (192.168/16 prefix)

Important: If you think you might want to connect to the Internet in the future, you should register with an Internet registry and use the IP addresses provided by the registry when setting up your private network. Otherwise, when you do connect to the Internet, you'll need to reconfigure every computer on your network.

If you set up a private TCP/IP network, you can also provide DNS service. By setting up TCP/IP and DNS on your local area network, your users will be able to easily access file, web, mail, and other services on your network.

Hosting Several Internet Services With a Single IP Address

You may have one server supplying all your Internet services (like mail, web). They may all be running on one computer with a single IP address. For example, you may want to have the domain name `www.example.com` resolve to the same IP address as `ftp.example.com`, or `mail.example.com`. This appears to be several servers to anyone accessing the services, but they are all really one server at one IP address.

Setting up the DNS records for this service is easy. You'll just add aliases to the machine DNS record. Setting up the DNS names for these services does not enable or configure the services, it merely makes easy to remember names for each service offered. This can ease setup and configuration of the client software for each service.

For example, for every service you want to show:

- Create `mail.example.com` to enter on mail clients.
Make sure to check the mail server box on the machine panel.
- Create `www.example.com` to enter on web browsers.
- Create `afp.example.com` for Apple File Sharing in the Finder.
- Create `ftp.example.com` to enter in ftp clients.

As your needs grow, you can add other computers to the network to take over these services. Then all you have to do is remove the alias from the machine's DNS record, and create a new record for the new machine, and your client's settings can remain the same.

Hosting Multiple Domains on the Same Server

You may have one server supplying all your Internet services (like mail, web) for several different domain names. For example, you may need to have the domain name `www.example.com` resolve to the same IP address as `www.example.org`. This appears to be several servers to anyone accessing the domain, but they are all really one server at one IP address.

Setting up the DNS records for this service is easy. You'll just add aliases of the other domain names to the main server's machine DNS record panel. Setting up the DNS names for these services does not enable or configure the service for these domain names. This is used in conjunction with virtual domain hosting in mail and web services.

Where to Find More Information

For more information on DNS and BIND, see the following:

- *DNS and BIND, 4th edition*, by Paul Albitz and Cricket Liu (O'Reilly and Associates, 2001)
- The International Software Consortium website:
www.isc.org and www.isc.org/products/BIND/
- The DNS Resources Directory:
www.dns.net/dnsrd/

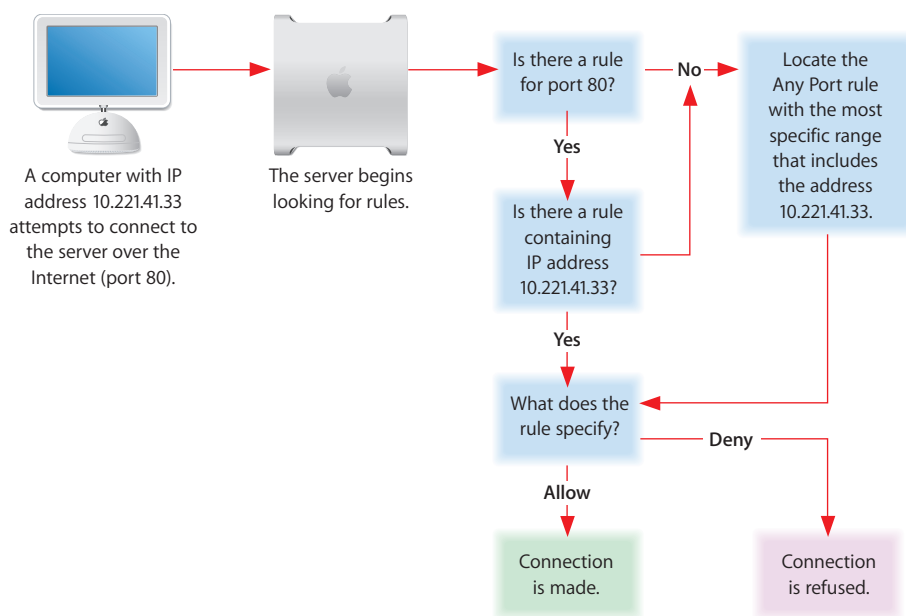
Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at the website www.ietf.org/rfc.html.

- A, PTR, CNAME, MX -For more information, see RFC 1035
- AAAA- For more information, see RFC 1886.

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. Firewall service scans incoming IP packets and rejects or accepts these packets based on the set of rules you create. You can restrict access to any IP service running on the server, and you can customize rules for all incoming clients or for a range of client IP addresses.

The illustration below shows an example firewall process.



Services such as Web and FTP are identified on your server by a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port number. When a computer tries to connect to a service, firewall service scans the rule list for a matching port number.

When a packet arrives at a network interface, and the firewall is enabled, the packet is compared against each rule, starting with the lowest-numbered (highest-priority) rule. When a rule matches the packet, the action specified in the rule (such as allow or deny) is taken. Then, depending on the action, additional rules may be checked.

The rules you create are applied to TCP packets and can also be applied to UDP packets. In addition, you can set up rules for restricting Internet Control Message Protocol (ICMP), or Internet Group Management Protocol (IGMP) using advanced rule creation.

Important: When you start firewall service the first time, only ports essential to remote administration of the server are open, including Secure Shell (22), and several others. Additional ports are dynamically opened to allow specific responses to queries initiated from the server. If you want to allow remote access to other services on your computer, you need to open additional ports, and you can do so using the Services section of the Settings panel.

If you plan to share data over the Internet, and you don't have a dedicated router or firewall to protect your data from unauthorized access, you should use firewall service. This service works well for small to medium businesses, schools, and small or home offices.

Large organizations with a firewall can use firewall service to exercise a finer degree of control over their servers. For example, individual workgroups within a large business, or schools within a school system, may want to use firewall service to control access to their own servers.

IP Firewall also provides stateful packet inspection which determines whether an incoming packet is a legitimate response to an outgoing request or part of an ongoing session, allowing packets that would otherwise be denied.

Basic Firewall Practices

By default, Mac OS X Server uses a simple model for a useful and secure firewall. If a firewall is too restrictive, the network behind it may as well be isolated. If a firewall is too permissive, it fails to secure the assets behind it from intrusion. Following the three aspects of the basic model will allow for maximum flexibility and utility with minimum unintended risk.

Allow essential IP activity

Essential IP activity includes those network activities necessary to use IP and function within an IP environment. These activities include operations such as loopback, and are expressed as high-priority (low-numbered) rules, visible in the Advanced panel of the firewall settings. These are automatically configured for you.

Allow service specific activity

Service specific activity refers to network packets destined for certain service-specific ports like web service, or mail service. By allowing traffic to ports with designated, configured services, you allow access through the firewall on a per-service basis. These are expressed as medium-priority rules, and correspond to check boxes in the Service panel of the firewall settings. You make these changes yourself based on your settings and address groups.

Deny all packets not already allowed

This is the final catch-all practice. If a packet or traffic to a port is unsolicited, that packet is discarded and not allowed to reach its destination. This is expressed as low-priority (high-numbered) rules, visible in the Advanced panel of the firewall settings. A basic set of “deny” rules for the firewall is created by default.

Firewall Startup

Although the firewall is treated as a service by the Server Admin application, it is not implemented by a running process like other services. It is simply a set of behaviors in the kernel, controlled by the `ipfw` and `sysctl` tools. To start and stop the firewall, the Server Admin application sets a switch using the `sysctl` tool. When the computer starts, a startup item named `IPFilter` checks the `/etc/hostconfig` file for the “`IPFILTER`” flag. If it is set, the `sysctl` tool is used to enable the firewall like so:

```
sysctl -w net.inet.ip.fw.enable=1
```

Otherwise, it disables the firewall like so:

```
sysctl -w net.inet.ip.fw.enable=0
```

Note that the rules loaded in the firewall remain there regardless of this setting. It's just that they are ignored when the firewall is disabled.

Like most startup items, the `IPFilter` startup item launches in a predetermined order, and only after certain prerequisite startup items have completed. In Mac OS X Server v10.4, the login window is presented while startup items may still be running. It is therefore possible to log in while the firewall has not been set to its configured settings. The startup item that sets up the firewall should generally finish within a few minutes of starting the system.

Understanding Firewall Rules

When you start firewall service, the default configuration denies access to all incoming packets from remote computers except ports for remote configuration. This provides a high level of security. Stateful rules are in place as well, so responses to outgoing queries initiated by your computer are also allowed. You can then add new IP rules to allow server access to those clients who require access to services.

To learn how IP rules work, read the following section. To learn how to create IP rules, see “Managing Firewall Service” on page 66.

What is a Firewall Rule?

A firewall rule is a set of characteristics of an IP packet, along with an action to be taken for each packet that matches the characteristics. The characteristics might include the source or destination address, source or destination port, protocol, or network interface. Addresses might be expressed as single IP address, or might include a range of addresses. A service port might be expressed as a single value, a list of values, or a range of values. The IP address and the subnet mask together determine the range of IP addresses to which the rule applies, and can be set to apply to all addresses.

IP Address

IP addresses consist of four segments with values between 0 and 255 (the range of an 8-bit number), separated by dots (for example, 192.168.12.12). The segments in IP addresses go from general to specific (for example, the first segment might belong to all the computers in a whole company, and the last segment might belong to a specific computer on one floor of a building).

Subnet Mask

A subnet mask indicates which segments in the specified IP address can vary on a given network and by how much. The subnet mask is given in Classless Inter Domain Routing (CIDR) notation. It consists of the IP address followed by a slash (/) and a number from 1 to 32, called the IP prefix. An IP prefix identifies the number of significant bits used to identify a network.

For example, 192.168.2.1 /16 means the first 16 bits (the first two numbers separated by periods) are used to represent the network (every machine on the network begins with 192.168) and the remaining 16 bits (the last two numbers separated by periods) are used to identify hosts (each machine has a unique set of trailing numbers).

Subnet masks can be given in another notation which is IP address followed by a colon (:) and the netmask. A netmask is a group of 4 numbers from 0-255, separated by periods which are the decimal equivalents to the CIDR notation's significant bits.

Addresses with subnet masks in CIDR notation correspond to address notation subnet masks.

CIDR	Corresponds to Netmask	Number of addresses in the range
/1	128.0.0.0	4.29x10 ⁹
/2	192.0.0.0	2.14x10 ⁹
/3	224.0.0.0	1.07x10 ⁹
/4	240.0.0.0	5.36x10 ⁸
/5	248.0.0.0	1.34x10 ⁸
/6	252.0.0.0	6.71x10 ⁷
/7	254.0.0.0	3.35x10 ⁷
/8	255.0.0.0	1.67x10 ⁷
/9	255.128.0.0	8.38x10 ⁶
/10	255.192.0.0	4.19x10 ⁶
/11	255.224.0.0	2.09x10 ⁶
/12	255.240.0.0	1.04x10 ⁶
/13	255.248.0.0	5.24x10 ⁵
/14	255.252.0.0	2.62x10 ⁵
/15	255.254.0.0	1.31x10 ⁵
/16	255.255.0.0	65536
/17	255.255.128.0	32768
/18	255.255.192.0	16384
/19	255.255.224.0	8192
/20	255.255.240.0	4096
/21	255.255.248.0	2048
/22	255.255.252.0	1024
/23	255.255.254.0	512
/24	255.255.255.0	256
/25	255.255.255.128	128
/26	255.255.255.192	64
/27	255.255.255.224	32
/28	255.255.255.240	16
/29	255.255.255.248	8
/30	255.255.255.252	4
/31	255.255.255.254	2
/32	255.255.255.255	1

Using Address Ranges

When you create an address group using Server Admin, you enter an IP address and a subnet mask. The three types of address notations allowed are:

- A single address: 192.168.2.1
- A range expressed with CIDR notation: 192.168.2.1/24
- A range expressed with netmask notation: 192.168.2.1:255.255.255.0

Server Admin shows you the resulting address range, and you can change the range by modifying the subnet mask. When you indicate a range of possible values for any segment of an address, that segment is called a *wildcard*. The following table gives examples of address ranges created to achieve specific goals.

Goal	Sample IP address	Enter this in the address field:	Address range affected
Create a rule that specifies a single IP address.	10.221.41.33	10.221.41.33 or 10.221.41.33/32	10.221.41.33 (single address)
Create a rule that leaves the fourth segment as a wildcard.	10.221.41.33	10.221.41.33/24	10.221.41.0 to 10.221.41.255
Create a rule that leaves part of the third segment and all of the fourth segment as a wildcard.	10.221.41.33	10.221.41.33/22	10.221.40.0 to 10.221.43.255
Create a rule that applies to all incoming addresses.		Select "Any"	All IP addresses

Rule Mechanism and Precedence

The rules in the Settings > Services panel operate in conjunction with the rules shown in the Advanced Rules panel. Usually, the broad rules in the Advanced panel block access for all ports. These are lower-priority (higher-numbered) rules and take effect after the rules in the General panel. The rules created with the General panel open access to specific services, and are higher priority. They take precedence over those created in the Advanced panel. If you create multiple rules in the Advanced panel, a rule's precedence is determined by the rule number which is the rule's order in the Advanced panel. Rules in the advanced panel can be reordered by dragging the rule within the list.

For most normal uses, opening access to designated services in the advanced panel is sufficient. If necessary, you can add additional rules using the Advanced panel, creating and ordering them as needed.

Multiple IP Addresses

A server can support multiple homed IP addresses, but firewall service applies one set of rules to all server IP addresses. If you create multiple alias IP addresses, then the rules you create will apply to all of those IP addresses.

Setting Up Firewall Service for the First Time

Once you've decided which rules you need to create, follow these overview steps to set up firewall service. If you need more help to perform any of these steps, see "Managing Firewall Service" on page 66 and the other topics referred to in the steps.

Step 1: Learn and plan

If you're new to working with IP Firewall, learn and understand firewall concepts, tools, and features of Mac OS X Server and BIND. For more information, see "Understanding Firewall Rules" on page 62.

Then plan your IP Firewall Service by planning which services you want to provide access to. Mail, web, and FTP services generally require access from computers on the Internet. File and print services will most likely be restricted to your local subnet.

Once you decide which services you want to protect using firewall service, you need to determine which IP addresses you want to allow access to your server, and which IP addresses you want to deny access to your server. Then you can create the appropriate rules.

Step 2: Start firewall service

In Server Admin, select Firewall and click Start Service. By default, this blocks all incoming ports except those used to configure the server remotely. If you're configuring the server locally, turn off external access immediately.

Important: If you add or change a rule after starting firewall service, the new rule will affect connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server will be disconnected.

Step 3: Create an IP address group that rules will apply to

By default, there is an IP address group created for all incoming IP addresses. Rules applied to this group will effect all incoming network traffic.

See "Creating an Address Group" on page 67 for more information.

Step 4: Activate service rules for each address group

In the Services panel, you can activate rules based on address groups as destination IP numbers.

For information about activating service rules, see "Opening the Firewall for Standard Services" on page 68.

Step 5: Create advanced rules (optional)

Read "Understanding Firewall Rules" on page 62 to learn how IP rules work. You use this to further configure all other services, strengthen your network security, and fine-tune your network traffic through the firewall.

By default, all UDP are blocked, except those in response to an outgoing query. You should apply rules to UDP ports sparingly, if at all, because denying certain UDP responses could inhibit normal networking operations.

If you rule UDP ports, don't select the "Log all allowed packets" option in the rule configuration windows in Server Admin. Since UDP is a "connectionless" protocol, every packet to a UDP port will be logged if you select this option.

For information about creating a new rule, see "Creating an Advanced IP Firewall Rule" on page 70.

Step 6: Save firewall service changes

Once you have configured your rules and determined which services to allow, save your changes so the new settings take effect.

Important: If you add or change a rule after starting firewall service, the new rule will affect connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server will be disconnected.

Managing Firewall Service

This section gives step-by-step instructions for starting, stopping, and configuring firewall address groups and rules.

Managing Panther Server 10.3 Firewalls with Tiger Server 10.4 Server Admin

Panther Server 10.3 does not support adding to the standard port rules, or drag-and-drop arrangement of rules.

If you are administering a Panther10.3 server firewall with a Tiger Server 10.4's Server Admin, you will not be able to edit the standard port rules or rearrange the rules. You will not have access to those aspects of Server Admin when connected to a Panther 10.3 server.

Starting and Stopping Firewall Service

By default, firewall service blocks all incoming TCP connections and denies all UDP packets, except those in response to outgoing requests from the server. Before you turn on firewall service, make sure you've set up rules allowing access from IP addresses you choose. Otherwise, no one will have access to your server.

Important: If you add or change a rule after starting firewall service, the new rule will affect connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server will be disconnected.

To start or stop firewall service:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Start Firewall.

When the service is started, the Stop Service button is available.

Creating an Address Group

You can define groups of IP addresses for your firewall rules. These groups are used to organize and target the rules. The “any” address group is for all addresses. Two other IP address groups are present by default, intended for the entire “10-net” range of private addresses, and the entire “192.168-net” range of private addresses.

Addresses can be listed as individual addresses (192.168.2.2), IP address and CIDR notation (192.168.2.0/24), or IP address and netmask notation (192.168.2.0:255.255.255.0).

To create an address group:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Click the Add (+) button to the right the Address Group pane.
- 5 Enter a group name.
- 6 Enter the addresses and subnet mask you want the rules to effect.
Use the Add (+) and Delete (-) buttons.
Use the word “any” to indicate any IP address.
- 7 Click OK.
- 8 Click Save.

Editing or Deleting an Address Group

You can edit your address groups to change the range of IP addresses effected. The default address group is for all addresses. You can remove address groups from your firewall rule list. The rules associated with those addresses are also deleted.

Addresses can be listed as individual addresses (192.168.2.2) or IP address and CIDR format netmask (192.168.2.0/24).

To edit or delete an address group:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the group name from the Address Group pane.

- 5 Click the Edit (/) button to the right the Address Group pane to edit it.
Click the Delete (-) button to the right the Address Group pane to delete it.
- 6 Edit the Group name or addresses as needed, and click OK.
- 7 Click Save.

Duplicating an Address Group

You can duplicate address groups from your firewall rule list. This can help speed up configuration of similar address groups.

To duplicate an address group:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the group name from the Address Group pane.
- 5 Click the Duplicate button to the right the Address Group pane.

Opening the Firewall for Standard Services

By default, firewall service blocks incoming TCP connections on ports that are not essential for remote administration of the server, and allows all UDP connections. Also, by default, stateful rules are in place that allow specific responses to outgoing requests. Before you turn on firewall service, make sure you've set up rules allowing access from IP addresses you choose; otherwise, no one will have access to your server.

You can easily allow standard services through the firewall without advanced and extensive configuration. Standard services include (but are not limited to):

- SSH access
- Web service
- Apple File service
- Windows File service
- FTP service
- Printer Sharing
- DNS/Multicast DNS
- ICMP Echo Reply (incoming pings)
- IGMP (Internet Gateway Multicast Protocol)
- PPTP VPN
- L2TP VPN
- QTSS media streaming
- iTunes Music Sharing

Important: If you add or change a rule after starting firewall service, the new rule will affect connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server will be disconnected.

To open the firewall for standard services:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Services tab.
- 4 Select an address group from the Edit Services pop-up menu.
- 5 Choose either to allow all traffic for the address group, or to allow traffic on designated points.
- 6 Check Allow for each services you want to allow to the address group.
If you don't see the service you need, you can add a port and description to the services list.
If you want to create a custom rule, see "Creating an Advanced IP Firewall Rule" on page 70.
- 7 Click Save.

Adding to the Services List

You can add custom ports to the services list. This will allow you to open specific ports to your address groups without having to create an advanced IP rule.

To add to the services list:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Services tab.
- 4 Click the Add (+) button below the services list.
- 5 Enter a rule name for the service.
- 6 Enter a single port (for example, 22) or a port range (for example, 650-750).
- 7 Choose a protocol.
If you want a protocol other than TCP or UDP, you need to use the Advanced panel to create a custom rule.
- 8 Click OK
- 9 Click Save.

Editing or Deleting Items in the Services List

You can remove or edit the ports to the services list. This will allow you to customize your services choices for ease of configuration.

To change the services list:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Services tab.
- 4 Select the service you want to change.
- 5 Click the Edit (/) button below the services list to edit it.
Click the Delete (-) button below the services list to delete it.
- 6 Edit the name, port, or protocol as needed, and click OK.
- 7 Click Save.

Creating an Advanced IP Firewall Rule

You can use the Advanced Settings pane to configure very specific rules for IP Firewall. IP firewall rules contain originating and destination an IP addresses with subnet masks. They also specify what to do with the network traffic received. You can apply a rule to all IP addresses, a specific IP address, or a range of IP addresses.

Addresses can be listed as individual addresses (192.168.2.2) or ranges defined by an IP address and CIDR netmask (192.168.2.0/24).

To create an advanced IP firewall rule:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Advanced Rules tab.
- 4 Click the Add (+) button.
Alternatively, you can select a rule similar to the one you want to create, and click Duplicate then Edit.
- 5 Select whether this rule will allow or deny access in the Action pop-up menu.
If you choose Other, enter the action desired (for example, log).
- 6 Choose a from the Protocol pop-up menu.
If you choose Other, enter the protocol desired (for example, icmp, esp, ipencap).
- 7 Choose a service from the pop-up menu.
If you want to select a nonstandard service port, choose Other.
- 8 If desired, choose to log packets that match the rule.

- 9 Choose an address group from the pop-up menu as the source of filtered traffic.
If you don't want to use an existing address group, enter the source IP address range (with CIDR notation) you want to filter.
If you want it to apply to any address, choose "any" from the pop-up menu.
- 10 If you have selected a nonstandard service port, enter the source port number.
- 11 Choose an address group from the pop-up menu as the destination of filtered traffic.
If you don't want to use an existing address group, enter the destination IP address range (with CIDR notation).
If you want it to apply to any address, choose "any" from the pop-up menu.
- 12 If you have selected a nonstandard service port, enter the destination port number.
- 13 Choose which network interface this rule applies to.
"In" refers to the designated WAN interface.
"Out" refers to the designated LAN interface.
If you select Other, enter the interface name (en0, en1, fw1, and so on)
- 14 Click OK.
- 15 Click Save to apply the rule immediately.

Editing or Deleting Advanced IP Firewall Rules

You can remove or edit advanced IP firewall rules. If you only want to disable a rule, and you might use it again, you could deselect the rule rather than deleting it.

If you edit a rule after turning on firewall service, your changes affect connections already established with the server. For example, if any computers are connected to your Web server, and you change the rule to deny all access to the server, connected computers will be disconnected.

To change an advanced IP firewall rule:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Advanced Rules tab.
- 4 Select the rule you want to change.
- 5 Click the Edit (/) button below the services list to edit it.
Click the Delete (-) button below the services list to delete it. If you're only deleting a rule, you've finished.
- 6 Edit the rule as needed, and click OK.
- 7 Click Save.

Changing the Advanced IP Firewall Rule Order

The order of the advanced IP firewall rules is determined by their order in the Advanced Rules tab.

To change the rule order:

- Drag the rules into the desired order.

Enabling Stealth Mode

You can hide the existence of your firewall by choosing not to send a connection failure notification to any connection that is blocked by the firewall. This effectively “hides” your server’s closed ports. For example, if a network intruder tries to connect to your server, even if the port is blocked, he knows that there is a server and may find other ways to intrude. If “stealth mode” is enabled, instead of being rejected, he won’t receive any indication that an attempted connection ever took place.

To enable stealth mode:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Advanced Rules tab.
- 4 Select “Enable for TCP” and/or “Enable for UDP” as desired.
- 5 Click Save.

Resetting an Unreachable Server

Due to some error with the firewall configuration, a server may become unreachable for remote administration. In such a case, the firewall will need to be reset to its default state, so Server Admin can administer the server.

This recovery procedure must be done by an administrator who has physical access to the server. This procedure requires use of, and facility with, the command-line interface.

To reset the firewall to its default:

- 1 Disconnect the server from the external Internet.
- 2 Restart the server in single-user mode by holding down the Command-s keys during start-up.
- 3 Remove or rename the address groups file.
This is found at `/etc/ipfilter/ip_address_groups.conf`.
- 4 Remove or rename the ipfw configuration file.
This is found at `/etc/ipfilter/ipfw.conf`.
- 5 Force flush the firewall rules by entering:

```
ipfw -f flush
```
- 6 Edit `/etc/hostconfig` and set `IPFILTER=-YES-`.

- 7 Finish starting up Mac OS X Server to the login window by typing:

```
exit
```

The machine will boot with the default firewall rules and with the firewall enabled, and Server Admin can then be used to refine the firewall configuration.

- 8 Log in to your server's local administrator account to confirm that the firewall is restored to its default configuration.
- 9 Reconnect your host to the Internet.

Monitoring Firewall Service

Firewalls are a network's first line of defense against malicious computer users (commonly called "hackers"). To maintain the security of your computers and users, you need to monitor firewall activity and deter potential threats. This section explains how to log and monitor your firewall.

Understanding the Active Rules Panel

The Active rule Panel shows counts of packets and bytes associated with each rule. When a change is made to the configuration of the firewall using Server Admin, the old firewall rules are flushed, new ones are generated and saved in a file, and the `ipfw(1)` command is invoked to load the rules into service. As part of the flush operation, the counts of packets and bytes associated with each rule are cleared.

The Active Rules panel reflects a snapshot of the state of the firewall at a point in time. When viewing this panel, note that there may be dynamic rules shown along with the static rules. These dynamic come and go in a matter of seconds, in response to network activity. They are the result of stateful rules (rules that include a "keep-state" clause). The Active Rules panel shows the rule number of the stateful rule that was triggered to create the dynamic rule.

Viewing the Firewall Status Overview

The Status Overview shows a simple summary of the firewall service. It shows the number of active rules, whether the service is running, and how many packets have been handled by the firewall.

To see the overview:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click the Overview button.

Viewing the Active Firewall Rules

The Active Rules panel shows a simple summary of the firewall rules. It shows:

- The rules in ipfw code format
- Each rule's priority
- Each rule's packet count
- Each rule's total of bytes handled

To see the overview:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click the Active Rules button.

Setting Up Logs for Firewall Service

You can log only the packets that are denied by the rules you set, only the packets that are allowed, or both. Both logging options can generate a lot of log entries, but there are ways to limit the volume:

- Log only the allowed packets, or only the denied packets, instead of all packets.
- Log only packets as long as necessary.
- Limit the total number of packets using the Logging Settings panel.
- Add a "count" rule in the Advanced Settings panel to tally the number of packets that match the characteristics you're interested in measuring.

You can choose to log allowed packets, denied packets, and a designated number of packets.

To set up logs:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the logging options you want.
- 5 Click Save to start logging.

Viewing the Firewall Log

Each rule you create in Server Admin corresponds to one or more rules in the underlying firewall software. Log entries show you the rule applied, the IP address of the client and server, and other information.

The log view is the contents of `/var/log/ipfw.log`. You can further filter the rules with the text filter box.

To view the log for firewall service:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Log tab.

Here are some examples of firewall log entries and how to read them.

Log Example 1

```
Dec 12 13:08:16 ballch5 mach_kernel: ipfw: 65000 Unreach TCP
    10.221.41.33:2190 192.168.12.12:80 in via en0
```

This entry shows that firewall service used rule 65000 to deny (unreach) the remote client at 10.221.41.33:2190 from accessing server 192.168.12.12 on Web port 80 via Ethernet port 0.

Log Example 2

```
Dec 12 13:20:15 mayalu6 mach_kernel: ipfw: 100 Accept TCP
    10.221.41.33:721 192.168.12.12:515 in via en0
```

This entry shows that firewall service used rule 100 to allow the remote client at 10.221.41.33:721 to access the server 192.168.12.12 on the LPR printing port 515 via Ethernet port 0.

Log Example 3

```
Dec 12 13:33:15 smithy2 mach_kernel: ipfw: 10 Accept TCP
    192.168.12.12:49152 192.168.12.12:660 out via lo0
```

This entry shows the NAT divert rule, applied to an outbound packet. In this case it diverts the rule to service port 660, which is the port the NAT daemon uses.

Viewing Denied Packets

Viewing denied packets can help you identify problems and troubleshoot firewall service.

To view denied packets:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Make sure “Log denied packets” is checked.
- 5 View log entries in Server Admin by clicking the Log button.
- 6 Enter the word “unreach” in the text filter box.

Viewing Packets Logged by Firewall Rules

Viewing the packets filtered by the firewall rules can help you identify problems and troubleshoot firewall service.

To view filtered packets:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Make sure “Log all allowed packets” is checked.

See “Editing or Deleting Advanced IP Firewall Rules” on page 71 if you have not turned on logging for a particular rule.

- 5 View log entries in Server Admin by clicking the Log button.
- 6 Enter the word “Accept” in the text filter box.

Troubleshooting Advanced IP Firewall Rules

The Firewall configuration Advanced panel will accept any input, assuming you are correctly configuring a rule. Errors are not noticed until the rules are saved and Server Admin applies all the rules using the `ipfw` command. Then, the first rule with a syntax error will cause the operation to stop, and an error message will appear in the log. This error will not indicate which rule is invalid, but all the valid rules before the invalid one will be loaded into the firewall. Here is the technique for figuring out which rule is invalid.

To figure out which rule is invalid:

- 1 Note the message in the log.
- 2 Wait a few minutes for Server Admin to display the active rules in the Overview section.
- 3 Compare the list of active rules in the Overview with the rule list in the Settings section.
- 4 Inspect the contents of `/etc/ipfilter/ipfw.conf.apple` file to see which ones Server Admin tried to load into the firewall.

The first one in that file that is not present in the Overview panel is almost surely the invalid one. There may also be additional invalid rules after that one.

- 5 If the rule corresponds to one from the advanced panel, you can disable it or correct it. Disabled rules appear in that `/etc/ipfilter/ipfw.conf.apple` preceded by a comment character so they are not processed by the `ipfw` tool.

Practical Examples

The IP firewall rules you create work together to provide security for your network. The examples that follow show how to use rules to achieve some specific goals.

Using IP Firewall with NAT

IP Firewall must be enabled to use NAT (Network Address Translation). Enabling NAT automatically creates a divert rule in the Firewall configuration. Although the Server Admin application in Tiger Server allows the NAT service and the Firewall service to be enabled and disabled independently, for the NAT service to operate, both the NAT and the Firewall service need to be enabled. An essential part of NAT is the packet divert rule used in the Firewall.

Warning: IP Firewall must be enabled for NAT to function.

The IP Firewall rule created tells the firewall how to route network traffic coming from the network behind the NAT gateway. When you have a LAN behind a NAT gateway, you need to create or be aware of the address group that corresponds to the LAN.

The easiest way to configure IP Firewall to work with NAT is to use the Gateway Setup Assistant. It will automatically configure the IP address groups in the Firewall, as well as create the proper packet divert rule. If you are setting up a NAT LAN through a gateway for the first time, Apple recommends that you use the Gateway Setup Assistant.

If you do not want to use the Gateway Setup Assistant, or have existing gateway settings you do not want overwritten, you can configure NAT and the IP Firewall manually.

For detailed instructions on setting up a NAT LAN, see “Linking a LAN to the Internet Through One IP Address” on page 93.

Block Web Access to Internet Users

This section shows you, as an example, how to allow users on your subnet access to your server’s Web service, but deny access to the general public on the Internet.

For this example, your local network has the private IP address range of 10.0.1.1 to 10.0.1.254. Your server’s web service is at 10.0.2.1 on the server’s en2 port.

To do this via an advanced rule:

- 1 In Server Admin, create an address group called “LAN” with the address range 10.0.1.1/24

This includes all addresses in the 10.0.1.x subnet range.

See “Creating an Address Group” on page 67 for instructions.

- 2 Create an advanced rule with the following settings:

- Action: Allow
- Protocol: TCP
- Service: Web
- Source address group: LAN
- Destination address: Other 10.0.2.1
- Interface: en2

See “Creating an Advanced IP Firewall Rule” on page 70 for instructions.

To do this via the standard rules:

- 1 In Server Admin, create an address group “Web Server” with the address range 10.0.2.1

See “Creating an Address Group” on page 67 for instructions.

- 2 Click Settings.
- 3 Select the Services tab.
- 4 Select the address group “Web Server” from the Edit Services pop-up menu.
- 5 Choose to allow traffic for the group “Web Server” on the designated web services port. Select Allow Web Service.
- 6 Click Save.

Logging Internet Access by Local Network Users

This section shows you, as an example, how to allow users on your LAN access to other servers’ Web service, and log their access to the general public on the Internet:

For this example, your local network has the private IP address range of 10.0.1.1 to 10.0.1.254.

- 1 In Server Admin’s IP Firewall panel, click Settings.
- 2 Select the Services tab.
- 3 Select the address group “any” from the Edit Services pop-up menu.
- 4 Choose to allow traffic for the group “Web Server” on the designated web services port. Select Allow Web Service.
- 5 Click Save.

- 6 Click the General tab.
- 7 Select Log All Allowed Packets.
View the logs in the Log panel.

Block Junk Mail

This section shows you, as an example, how to reject email from a junk mail sender with an IP address of 17.128.100.0 and accept all other Internet email.

Important: Set up very specific address ranges in rules you create to block incoming SMTP mail. For example, if you set a rule on port 25 to deny mail from all addresses, you'll prevent any mail from being delivered to your users.

To do this:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Services tab.
- 4 Select the "any" address group in the pop-up menu.
- 5 Enable "SMTP Mail."
- 6 Select the General tab.
- 7 Click the Add (+) button to create an address range.
- 8 Name the address group.
- 9 Enter 17.128.100.0 to the address range to indicate the junk mail sender's address.
- 10 Click OK.
- 11 Select your newly created address group.
- 12 Deselect "SMTP Mail" in the Services tab to disable mail transfer.
- 13 Click Save.

Allow a Customer to Access the Apple File Server

This section shows you, as an example, how to allow a customer with an IP address of 10.221.41.33 to access an Apple file server.

To do this:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Services tab.
- 4 Select the "any" address group.
- 5 Disable "Apple File Service" in the service pane.
- 6 Select the General tab.

- 7 Click the Add (+) button to create an address range.
- 8 Name the address group.
- 9 Enter 10.221.41.33 to the address range to indicate the customer's address.
- 10 Click OK.
- 11 Select the Services tab.
- 12 Select your newly created address group.
- 13 Select "Apple File Service" in the service pane to enable file access.
- 14 Click Save.

Common Network Administration Tasks That Use Firewall Service

Your firewall is the first line of defense against unauthorized network intruders, malicious users, and network virus attacks. There are many ways that such attacks can harm your data or use your network resources. This section lists a few of the common uses of firewall service in network administration.

Preventing Denial of Service (DoS) Attacks

When the server receives a TCP connection request from a client to whom access is denied, by default it sends a reply rejecting the connection. This stops the denied client from resending over and over again. However, a malicious user can generate a series of TCP connection requests from a denied IP address and force the server to keep replying, locking out others trying to connect to the server. This is one type of denial of service attack.

To prevent ping denial of service attacks:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Settings tab.
- 4 Select the "any" address group.
- 5 Deselect "ICMP Echo (ping) reply."
- 6 Click Save.

Important: Denial of service attacks are somewhat rare, so make these settings only if you think your server may be vulnerable to an attack. If you deny ICMP echo replies, services that use ping to locate network services will be unable to detect your server.

Controlling or Enabling Peer-to-Peer Network Usage

Sometimes network administrators need to control the use of Peer-to-Peer (P2P) file sharing applications. Such applications might use network bandwidth and resources inappropriately or disproportionately. P2P file sharing might also pose a security or intellectual property risk for a business.

You can cut off P2P networking by blocking all traffic incoming and outgoing on the port number used by the P2P application. You'll have to determine the port used for each P2P network in question. By default, Mac OS X Server's firewall blocks all ports not specifically opened.

You can choose to limit P2P network usage to IP addresses behind the firewall. To do so, you'll need to open the P2P port for your LAN interface, but continue to block the port on the interface connected to the Internet (WAN interface). To learn how to make a firewall rule, see "Creating an Advanced IP Firewall Rule" on page 70.

Controlling or Enabling Network Game Usage

Sometimes network administrators need to control the use of network games. The games might use network bandwidth and resources inappropriately or disproportionately.

You can cut off network gaming by blocking all traffic incoming and outgoing on the port number used by the game. You'll have to determine the port used for each network game in question. By default, Mac OS X Server's firewall blocks all ports not specifically opened.

You can choose to limit network game usage to IP addresses behind the firewall. To do so, you'll need to open the appropriate port on your LAN interface, but continue to block the port on the interface connected to the Internet (WAN interface). Some games require a connection to a gaming service for play, so this may not be effective. To learn how to make a firewall rule, see "Creating an Advanced IP Firewall Rule" on page 70.

You can open the firewall to certain games, allowing network games to connect to other players and game services outside the firewall. To do this, you'll need to open up the appropriate port on your LAN and WAN interface. Some games require more than one port to be open. Consult the game's documentation for networking details. To learn how to make a firewall rule, see "Creating an Advanced IP Firewall Rule" on page 70.

Port Reference

The following tables show the TCP and UDP port numbers commonly used by Mac OS X computers and Mac OS X Servers. These ports can be used when you're setting up your rules. See the following website to view the RFCs referenced in the tables:

www.faqs.org/rfcs

TCP port	Used for	Reference
7	echo	RFC 792
20	FTP data	RFC 959
21	FTP control	RFC 959
22	SSH (secure shell) Open Directory replica setup	
23	Telnet	RFC 854
25	SMTP (email)	RFC 821
53	DNS	RFC 1034
79	Finger	RFC 1288
80	HTTP (Web)	RFC 2068
88	Kerberos V5 KDC	RFC 1510
106	Open Directory Password Server (along with 3659)	
110	POP3 (email)	RFC 1081
111	Remote Procedure Call (RPC)	RFC 1057
113	AUTH	RFC 931
115	sftp	
119	NNTP (news)	RFC 977
123	Network Time Server synchronization (NTP)	RFC 1305
137	Windows Names	
138	Windows Browser	
139	Windows file and print service (SMB/CIFS)	RFC 100
143	IMAP (email access)	RFC 2060
201-208	AppleTalk	
311	Server Admin SSL, AppleShare IP remote Web administration, Server Monitor, Server Admin (servermgrd), Workgroup Manager (DirectoryService)	

TCP port	Used for	Reference
389	LDAP (directory) Sherlock 2 LDAP search	RFC 2251
407	Timbuktu	
427	SLP (service location)	
443	SSL (HTTPS)	
445	Microsoft Domain Server	
497	Dantz Retrospect	
514	shell, syslog	
515	LPR (print spooling)	RFC 1179
532	netnews	
548	AFP (Apple File Service)	
554	Real-Time Streaming Protocol (QTSS)	RFC 2326
591	FileMaker web access	
600–1023	Mac OS X RPC-based services (for example, NetInfo)	
625	Remote Directory Access	
626	IMAP Administration (Mac OS X mail service and AppleShare IP 6.x mail)	
631	IPP (printer sharing)	
636	LDAP SSL	
660	Server Settings, Server Manager	
687	AppleShare IP Shared Users and Groups, Server Monitor, Server Admin (servermgrd)	
749	Kerberos administration and changepw using the kadmind command-line tool	
985	NetInfo static port	
993	IMAP over SSL (mail)	
995	POP3 over SSL (mail)	
1085	Web Objects	
1099, 8043	Remote RMI and RMI/IIOP access to JBoss	
1220	QTSS Admin	
1694	IP Failover	
1723	PPTP VPN	RFC 2637
2049	NFS	

TCP port	Used for	Reference
2236	Macintosh Manager	
2399	FileMaker data access layer	
3004	iSync	
3031	Program Linking, Remote AppleEvents	
3283	ARD 2.0	
3306	MySQL	
3632	XCode distributed compiler	
3659	Open Directory Password Server (along with 106)	
3689	iTunes music sharing	
4111	XGrid	
5003	FileMaker name binding and transport	
5100	Camera and scanner sharing	
5190	iChat, and iChat file transfer	
5222	iChat server	
5223	iChat server SSL	
5269	iChat server -server to server	
5298	iChat -local subnet	
5432	ARD 2.0 database	
5900	ARD 2.0 VNC	
7070	Real-Time Streaming Protocol (QTSS)	
7777	iChat Server- file transfer proxy	
8000–8999	Web service	
8000-8001	QTSS MP3 streaming	
8005	Tomcat remote shutdown	
8043, 1099	Remote RMI and RMI/IIOP access to JBoss	
8080, 8443, 9006	Tomcat standalone and JBoss	
8080	Web service alternate (Apache 2 default)	
9007	Remote web server access to AIP port	
16080	Web service with performance cache redirect	
42000-42999	iTunes radio streams	

UDP port	Used for	Reference
7	echo	
53	DNS	
67	DHCP server (BootP), NetBoot server	
68	DHCP client	
69	Trivial File Transfer Protocol (TFTP)	
111	Remote Procedure Call (RPC)	
123	Network Time Protocol	RFC 1305
137	Windows Name Service (WINS)	
138	Windows Datagram Service (NETBIOS)	
161	Simple Network Management Protocol (SNMP)	
192	AirPort administration	
427	SLP (service location)	
497	Retrospect	
500	VPN ISAKMP/IKE	
513	who	
514	Syslog	
554	Real-Time Streaming Protocol (RTSP)	
600–1023	Mac OS X RPC-based services (for example, NetInfo)	
626	Serial number support	
985	NetInfo (when a shared domain is created using NetInfo Domain Setup)	
1701	VPN L2TP	
3283	ARD 1.2	
5353	Multicast DNS (mDNSResponder)	
2049	Network File System (NFS)	
3031	Program Linking	
3283	Apple Network Assistant, Apple Remote Desktop	
4500	IKE NAT traversal	
5060	iChat initiation	
5297, 5678	iChat - local	

UDP port	Used for	Reference
5353	Multicast DNS (mDNSResponder)	
6970 -6999	QTSS RTP streaming	
7070	Real-Time Streaming Protocol alternate (QTSS)	
16384-16403	iChat audio/video RTP and RTCP	

Where to Find More Information

For more information about ipfw:

You can find more information about ipfw, the tool which controls IP firewall service, by accessing its man page. It explains how to access its features and implement them. To access the man page use the Terminal application to enter:

```
man ipfw
```

Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. The RFC section of the following website contains several RFC numbers for various protocols:
www.ietf.org/rfc.html

The IANA (Internet Assigned Number Authority) maintains a list of "well known ports," or TCP and UDP ports which have been assigned by the organization for various protocols. The list can be found at:
www.iana.org/assignments/port-numbers

Additionally, important multicast addresses are documented in the most recent Assigned Numbers RFC, currently RFC 1700.

Network Address Translation (NAT) is sometimes referred to as IP masquerading. NAT is used to allow multiple computers access to the Internet with only one assigned public or external IP address. NAT allows you to create a private network which accesses the Internet through a NAT router or gateway.

The NAT router takes all the traffic from your private network and remembers which internal address made the request. When the NAT router receives the response to the request, it forwards it to the originating computers. Traffic that originates from the Internet does not reach any of the computers behind the NAT router unless Port forwarding is enabled.

Using NAT with Other Network Services

Enabling NAT on Mac OS X Server often requires detailed control over DHCP, so DHCP is configured separately in Server Admin. To learn more about DHCP, see Chapter 2, “DHCP Service,” on page 23.

Enabling NAT also automatically creates a divert rule to the Firewall configuration. The Server Admin application in Mac OS X Server allows the NAT service and the Firewall service to be enabled and disabled independently. But for the NAT service to operate, both the NAT and the Firewall service need to be enabled. This is because an essential part of NAT is the packet divert rule. That rule is added to the Firewall when NAT service is enabled, but the Firewall service must be turned on for the packet divert rule, or any Firewall rule, to have any effect.

Warning: IP Firewall must be enabled for NAT to function.

NAT LAN Configuration Overview

In order to configure a network segment as a NAT LAN, you need to take several different steps. Each one is necessary to have a functioning private network behind a NAT gateway. A detailed example of setup can be found in “Linking a LAN to the Internet Through One IP Address” on page 93. The following section provides a high-level overview of the configuration process

Step 1: Choose your NAT gateway and its interface functions

It needs to be a Mac OS X Server computer with (at least) two network interfaces: one to connect to the Internet (the WAN port), and one to connect to your private network segment (the LAN port).

Step 2: Decide how the NAT LAN clients will get their IP addresses

You can assign your own static IP address within the approved ranges for private LANs, or you can use Mac OS X Server’s DHCP feature to assign addresses for you.

Step 3: Configure the gateway’s network settings

Assign you public IP address to the WAN port, and your internal gateway’s address to the LAN port.

Step 4: Configure NAT settings

See “Configuring NAT Service” on page 89.

Step 5: Configure port forwarding settings

See “Configuring Port Forwarding” on page 90.

Step 6: Start NAT service

See “Starting and Stopping NAT Service” on page 89.

Step 7: Start Firewall Service

For the NAT service to operate, both the NAT and the Firewall service need to be enabled. See “Starting and Stopping Firewall Service” on page 66 for more information.

Step 8: Configure and Start DHCP Service, if applicable

If the clients will have their addresses dynamically assigned, configure DHCP and start it now. See Chapter 2, “DHCP Service,” for more information.

Starting and Stopping NAT Service

You use Server Admin to start and stop NAT service on your default network interface. Starting NAT service does not automatically start DHCP on the NAT interface, so LAN addressing must be handled separately.

Starting the NAT service is not the same as configuring a network segment as a NAT LAN.

To start NAT service:

- 1 In Server Admin, select NAT from the Computers & Services pane.
- 2 Click Start Service.

When the service is running, Stop Service becomes available.

Configuring NAT Service

You use Server Admin to indicate which network interface is connected to the Internet or other external network.

Configuring the NAT service is not the same as configuring a network segment as a NAT LAN.

To configure NAT service:

- 1 In Server Admin, select NAT from the Computers & Services pane.
- 2 Click Settings.
- 3 Select “IP Forwarding and Network Address Translation.”
- 4 Choose the network interface from the “Network connection to share:” pop-up menu.
This interface should be the one that connects to the Internet or external network.
- 5 Click Save.

Creating a Gateway Without NAT

Sometimes you need to use a computer as a gateway between network segments, but you don't need to translate their IP addresses between public and private ranges. This is called "IP address forwarding." Mac OS X Server supports IP address forwarding through the NAT section of Server Admin.

For this configuration, you may have various network configurations. For example, some other server may be translating private IP address to public addresses using NAT, but your Mac OS X Server gateway may be routing information between various private address subnets. Likewise, you may want to run a firewall between network segments within your own LAN. Any condition in which you'd want to route network traffic through the server without masquerading IP addresses is a condition which involves IP address forwarding.

The steps for creating a gateway for address forwarding are the same as those for creating a NAT LAN. This means that the network ports must be configured to their proper settings, and the Firewall Service must be enabled for the gateway to function.

To configure a gateway without NAT service:

- 1 In Server Admin, select NAT from the Computers & Services pane.
- 2 Click Settings.
- 3 Select "IP Forwarding only."
- 4 Click Save.

Configuring Port Forwarding

You can direct incoming traffic to your NAT network to a specific IP address behind the NAT gateway. This allows you to set up computers on the internal network that handle certain incoming connections without exposing the other computers to outside connections. For example, you could set up a web server behind the NAT and forward all incoming TCP connection requests on port 80 to the designated web server.

You can't forward the same port to multiple computers, but you can forward any number of different ports to the same computer.

Enabling Port Forwarding requires use of, and facility with, the Terminal as well as administrative access to root privileges through `sudo`. You will need to edit a plist, and the contents of that plist will be used to generate `/etc/nat/natd.conf.apple`, which is passed to the NAT daemon when it is started. Do not try to edit `/etc/nat/natd.conf.apple` directly. If you choose to use a plist editor instead of a command line text editor, you may need to alter the following instructions to suit.

To forward port traffic:

- 1 If the file `/etc/natd.plist` doesn't exist, make a copy of the default NAT daemon plist.

```
sudo cp /etc/nat/natd.plist.default /etc/natd.plist
```

- 2 Using a Terminal editor, add a new block of XML text to `/etc/natd.plist` before the two last lines which end the file (`</dict>` and `</plist>`)

Add this block, and substitute your desired settings where indicated by italics:

```
<key>redirect_port</key>
<array>
<dict>
<key>proto</key>
<string>TCP or UDP</string>
<key>targetIP</key>
<string>LAN_ip</string>
<key>targetPortRange</key>
<string>LAN_ip_range</string>
<key>aliasIP</key>
<string>WAN_ip</string>
<key>aliasPortRange</key>
<string>WAN_port_range</string>
</dict>
</array>
```

- 3 Save your file changes.

The changes made, except for those settings that Server Admin can change and comments, will be respected by the server configuration tools (Server Admin, Gateway Setup Assistant and `serveradmin`).

- 4 Configure NAT service in Server Admin, as desired.

See “Configuring NAT Service” on page 89 for more information.

- 5 Click Save.

Port Forwarding Examples

You can forward a single port or any number of ports to a given IP address. The ports on the WAN side do not have to match the ports on the LAN side, but must correspond. For example, if you forward 10 consecutive ports from the WAN side, you must forward them to 10 consecutive ports on the LAN side, but they don't need to be the same 10.

Single Port Forwarding

This example shows the setting to forward TCP port 80 (web service) connections on the WAN address of 171.28.128.128 to TCP port 80 (web service) on the private LAN address of 192.168.1.1. The block of text to add to the `/etc/natd.plist` file is:

```
<key>redirect_port</key>
<array>
<dict>
<key>proto</key>
<string>TCP</string>
<key>targetIP</key>
<string>192.168.1.1</string>
<key>targetPortRange</key>
<string>80</string>
<key>aliasIP</key>
<string>17.128.128.128</string>
<key>aliasPortRange</key>
<string>80</string>
</dict>
</array>
```

Multiple Port Forwarding

This example shows the setting to forward TCP and UDP ports 600-1023 (NetInfo, full range) connections on the WAN address of 171.28.128.128 to corresponding ports on the private LAN address of 192.168.1.1. The block of text to add to the `/etc/natd.plist` file is:

```
<key>redirect_port</key>
<array>
<dict>
<key>proto</key>
<string>TCP</string>
<key>targetIP</key>
<string>192.168.1.1</string>
<key>targetPortRange</key>
<string>600-1023</string>
<key>aliasIP</key>
<string>17.128.128.128</string>
<key>aliasPortRange</key>
<string>600-1023</string>
</dict>
</array>
<array>
<dict>
<key>proto</key>
<string>UDP</string>
<key>targetIP</key>
<string>192.168.1.1</string>
<key>targetPortRange</key>
<string>600-1023</string>
<key>aliasIP</key>
<string>17.128.128.128</string>
```

```
<key>aliasPortRange</key>
<string>60-1023</string>
</dict>
</array>
```

Monitoring NAT Service

You might want to monitor your NAT service for troubleshooting and security. This section describes the NAT status overview and monitoring NAT divert activity.

Viewing the NAT Status Overview

The NAT status overview allows you to see if the service is running, and how many protocol links are active.

To see the overview:

- 1 In Server Admin, choose NAT Service from the Computers & Services list.
- 2 Click the Overview button.

Common Network Administration Tasks That Use NAT

The following sections illustrate some common network administration tasks that use NAT service.

Linking a LAN to the Internet Through One IP Address

The easiest way to link a NAT LAN to the Internet is to use the Gateway Setup Assistant. It will automatically configure the IP address groups in the Firewall, as well as create the proper packet divert rule. If you are setting up a NAT LAN through a gateway for the first time, it is recommended to use the Gateway Setup Assistant. See “Linking Your Network to the Internet” on page 15 for more information on the Gateway Setup Assistant.

If you do not want to use the Gateway Setup Assistant, or have existing gateway settings you do not want overwritten, you can configure NAT and the IP Firewall manually. To do so, you will need a Mac OS X Server with two network interfaces, one to connect to the Internet, and one to your private network. This example assumes the following configuration:

- *Ethernet Interface names and functions:* Ethernet Built-in (connected to Internet), PCI Ethernet Slot 1 (connected to internal network)
- *Internet or Public IP address:* 17.254.0.3 (example only, your IP number will be provided by your ISP)
- *Internet or Public DNS IP address:* 17.254.1.6 (example only, your IP number will be provided by your ISP)
- *Private Network IP address range and netmask:* 192.168.0.2–192.168.0.254 (also expressed as 192.168.0.0/24 or 192.168.0.0:255.255.255.0)
- *Server's Private Network IP address:* 192.168.0.1

- *LAN Client IP address settings*: Configure IPv4 Using DHCP
While not strictly necessary (NAT can be used with static IP addresses instead of DHCP), this allows easy configuration of client computers.

To configure your NAT LAN:

- 1 Open the Network pane of System Preferences on the gateway server.
- 2 In the active Network Port Configurations screen, make sure the interface “Ethernet Built-in” is on top of the list of interfaces.

If it isn’t, drag it to the top of the list. This sets the default gateway in the routing table. The top interface is always configured to be out the Internet or WAN.
- 3 Make sure the IP address and settings for “Ethernet Built-in” are your public address settings from your ISP. In this example they would be:
 - IP Address: 17.254.0.3
 - Netmask: 255.255.252.0
 - DNS: 17.254.1.6
- 4 Make sure the IP address and settings for “PCI Ethernet Slot 1” are your local address settings. In this example, they would be:
 - IP Address: 192.168.0.1
 - Netmask: 255.255.255.0
 - DNS: 17.254.1.6
- 5 Click Apply Changes, if necessary.
- 6 Open Server Admin.
- 7 In Server Admin, choose DHCP from the Computers & Services list.
- 8 In Server Admin, create an address group for the internal LAN with the following configuration parameters:
 - Subnet name: <whatever you want>
 - Starting IP Address: 192.168.0.2
 - Ending IP Address: 192.168.0.254
 - Subnet Mask: 255.255.255.0
 - Network Interface: en1
 - Router: 192.168.0.1
 - Lease time: <whatever you want>
 - DNS: 17.254.1.6See “Creating Subnets” on page 24 for detailed instructions for configuring DHCP.
- 9 Enable DHCP service.
- 10 In Server Admin, choose NAT from the Computers & Services list.
- 11 Configure NAT using the following setting:
 - Network Connection to Share: Built-In Ethernet
- 12 Click Save, if necessary.

- 13 Enable NAT Service.
- 14 In Server Admin, choose Firewall from the Computers & Services list.
- 15 Enable the Firewall.
- 16 Create Firewall rules to allow access to and from your private network.
For example, create an IP address group called “Private LAN” for the addresses 192.168.0.0/24.
See “Creating an Address Group” on page 67 for detailed instructions.
- 17 Enable any services you want the Private LAN to access (web, SSH, file sharing, and so on) using the “Private LAN” group.
See “Opening the Firewall for Standard Services” on page 68 for detailed instructions.
- 18 Enable any services you want the Internet to access on your private LAN (web, SSH, file sharing, and so on) using the “any” address group.
See “Opening the Firewall for Standard Services” on page 68 for detailed instructions.
- 19 Click Save.

Setting Up LAN Party For Gaming

Setting up a LAN party is essentially the same as “Linking a LAN to the Internet Through One IP Address.” Special considerations:

- Take special care to open the ports necessary to play an Internet-enabled game.
- If the game is to be played only within the LAN, you don’t have to open the Firewall to game ports.
- If you have computers joining and leaving the LAN, you’ll want to use DHCP to for client address configuration.

Setting Up “Virtual Servers”

A virtual server is a gateway server that sends services behind a NAT wall to real servers on a port-by-port basis. For example, you have a NAT gateway which is 17:100.0.1 (that is, domain.example.com) which could be set to forward web traffic (port 80) to 10.0.0.5 (port 80) behind the firewall, and requests for ssh traffic (port 22) could send the packets to 10.0.0.15 (port 22). In the example above, the NAT gateway is not really serving the web content, the server at 10.0.0.5 is, but that is invisible to the clients browsing the web site.

To the Internet you have one server, but behind the NAT barrier, you have as many or as few as you need. This can be used as load balancing, or as an organizational scheme for the network’s topography. Virtual servers also let you easily reroute network traffic to other computers on the LAN by reconfiguring the gateway.

Virtual servers require three service configurations: NAT, DNS, and IP Firewall.

The NAT service needs to be configured with port forwarding of the desired virtual port. The DNS record for the server should accept a few aliases of common services and resolve them all to the same IP address. Finally, the firewall needs to allow traffic on certain ports into the NAT LAN.

In this example, you'll set up a NAT gateway, and point two domain names and services to different computers behind the gateway firewall. Assume the following configuration details:

- *Ethernet Interface names and functions:* Ethernet Built-in (connected to Internet), PCI Ethernet Slot 1 (connected to internal network)
- *Internet or Public IP address:* 17.100.0.1 (example only, your IP number and netmask information will be provided by your ISP)
- *Private Network IP address range and netmask:* 192.168.0.0–192.168.0.255 (also expressed as 192.168.0.0/24 or 192.168.0.0:255.255.255.0)
- *Gateway Server's Private Network IP address:* 192.168.0.1
- *Web Server's Private Network IP address:* 192.168.0.2
- *Mail Server's Private Network IP address:* 192.168.0.3
- *Web and Mail Server's IP address settings:* Configure IPv4 Using DHCP
While not strictly necessary (NAT can be used with static IP addresses instead of DHCP), this allows easy configuration of client computers.

To configure your virtual servers

- 1 Open Server Admin.
- 2 In Server Admin, choose DHCP from the Computers & Services list.
- 3 In Server Admin, create an address group for the internal LAN with the following configuration parameters:
 - Subnet name: <whatever you want>
 - Starting IP Address: 192.168.0.2
 - Ending IP Address: 192.168.0.254
 - Subnet Mask: 255.255.255.0
 - Network Interface: en1
 - Router: 192.168.0.1
 - Lease time: <whatever you want>
 - DNS: <provided by ISP>
 - Static Mapping (web): <web server's Ethernet address> mapped to 192.168.0.2
 - Static Mapping (mail): <mail server's Ethernet address> mapped to 192.168.0.3See "Creating Subnets" on page 24 and "Assigning Static IP Addresses Using DHCP" on page 30 for detailed instructions for configuring DHCP.
- 4 Enable DHCP service.
- 5 In Server Admin, choose NAT from the Computers & Services list.
- 6 Configure NAT using the following setting:
 - Network Connection to Share: Built-In Ethernet

- Port forwarding: TCP port 80 (web) to 192.168.0.2
- Port forwarding: TCP port 25 (mail) to 192.168.0.3

7 Click Save.

8 Enable NAT Service.

9 In Server Admin, choose Firewall from the Computers & Services list.

10 Enable the Firewall.

11 Create Firewall rules to allow access to your private network.

See “Creating an Address Group” on page 67 for detailed instructions.

12 Enable the two services you want the Internet to access on your private LAN (web and SMTP mail) using the “any” address group.

See “Opening the Firewall for Standard Services” on page 68 for detailed instructions.

13 Click Save.

14 Add two aliases to your gateway server’s DNS record.

Contact your DNS provider (usually your ISP), and request an “A” record be added with the name “www.example.com” to the IP address 17100.0.1. Request an MX record with the name “mail.example.com” to the same IP address. These records are in addition to your existing A and CNAME records for your domain.

Now all web traffic to www.example.com will be forwarded to the internal server at 192.168.0.2, and incoming mail traffic to mail.example.com will be delivered to the internal server at 192.168.0.3.

If you want to change the servers behind the NAT (hardware upgrade, for example), now all you have to do is change the DHCP Static IP addressing to the Ethernet addresses of the new servers. The new servers will be assigned the existing internal IP addresses designated for web and mail, and the gateway will forward the traffic to the new servers seamlessly.

Where to Find More Information

For more information about natd:

You can find more information about natd, the daemon process which controls NAT service, by accessing its man page. It explains how to access its features and implement them. To access the man page use the Terminal application to enter:

```
man natd
```

Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at the website:

www.ietf.org/rfc.html

For NAT descriptions, see:

- RFC 1631
- RFC 3022.

Virtual Private Network (VPN) is two or more computers or networks (nodes) connected by a private link of encrypted data. This link simulates a local connection, as if the remote computer were attached to the local area network (LAN).

VPNs allow users at home or otherwise away from the LAN to securely connect to it using any network connection, such as the Internet. From the user's perspective, the VPN connection appears as a dedicated private link.

VPN technology also allows an organization to connect branch offices over the Internet, while maintaining secure communications. The VPN connection across the Internet acts as a wide area network (WAN) link between the sites.

VPNs have several advantages for organizations whose computer resources are physically separated. For example, each remote user or node uses the network resources of its Internet Service Provider (ISP) rather than having a direct, wired link to the main location. VPNs also allow verified mobile users to access private computer resources (file servers, and so on) from any connection to the Internet. Finally, VPN can link multiple LANs together over great distances using existing Internet infrastructure.

This chapter describes VPN authentication method, transport protocols, and how to configure, manage, and monitor VPN service. It does not include instructions for configuring VPN *clients* for use of your VPN server.

VPN and Security

VPNs stress security by strong authentication of identity, and encrypted data transport between the nodes, for data privacy and inalterability. The following section contains information about each supported transport and authentication method.

Transport Protocols

You'll be able to enable either or both of the encrypted transport protocols. Each has its own strengths and requirements.

Layer Two Tunneling Protocol, Secure Internet Protocol (L2TP/IPSec)

L2TP/IPSec uses strong IPsec encryption to “tunnel” data to and from the network nodes. It is based on Cisco’s L2F protocol. IPsec requires security certificates (either self-signed or from a Certificate Authority like Verisign), or a predefined shared secret between connecting nodes. The shared secret must be entered on the server as well as a client. It is not a password for authentication, nor does it generate encryption keys to establish secure tunnels between nodes. It is a token that allows the key management systems to trust each other.

L2TP is Mac OS X Server’s preferred VPN protocol due to its superior transport encryption and its ability to be authenticated via Kerberos.

Point to Point Tunneling Protocol (PPTP)

PPTP is a common VPN protocol as well as the Windows standard VPN protocol. PPTP offers good encryption (provided the passwords used are strong passwords) and supports a number of authentication schemes. It uses the user-provided password to produce an encryption key. You can also allow 40-bit (weak) security encryption in addition to the default 128-bit (stronger) encryption if needed by your VPN clients.

PPTP is necessary if you have old Windows clients or Mac OS X 10.2.x clients.

Authentication Method

Mac OS X Server L2TP VPN uses either Kerberos v5 or Microsoft’s Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) for authentication. Mac OS X Server PPTP VPN uses MS-CHAPv2, exclusively, for authentication.

Kerberos is a secure authentication protocol which depends on a Kerberos Key Distribution Server as a “trusted third party” to authenticate a client to a server. MS-CHAPv2 authentication doesn’t require the same authentication infrastructure as Kerberos. It method encodes passwords when they’re sent over the network, and stores them in a scrambled form on the server offering good security during network transmission. It is also the standard Windows authentication scheme for VPN.

Mac OS X Server PPTP VPN can use additional authentication methods. Each has its own strengths and requirements. It is not possible to choose any other authentication method for PPTP using Server Admin. If you need to configure a different authentication scheme from the default (for example, to use RSA Security’s SecurID authentication), you’ll need to edit the VPN configuration file manually. The configuration file is located at:

`/Library/Preferences/SystemConfiguration/com.apple.RemoteAccessServers.plist`

See “Offering SecurID Authentication With VPN Server” on page 107 for more information.

Before You Set Up VPN Service

Before setting up Virtual Private Network (VPN) service, you need to determine which transport protocol you're going to use. The table below shows which protocols are supported by different platforms.

If you have...	you can use L2TP/IPSec.	you can use PPTP.
Mac OS X 10.4 and 10.3.x clients	X	X
Mac OS X 10.2.x clients		X
Windows clients	X (if Windows XP)	X
Linux or Unix clients	X	X

If you're using L2TP, you need to have a Security Certificate (from a Certificate Authority or self-signed), or a predefined shared secret between connecting nodes. If you choose a shared secret, it needs to be secure as well (at least 8, but better yet 12 or more alphanumeric characters with punctuation and without spaces) and kept secret by the users.

If you're using PPTP, you need to make sure all of your clients support 128-bit PPTP connections, for greatest transport security. Be aware that enabling 40-bit transport security is a serious security risk.

Configuring other Network Services for VPN

Enabling VPN on Mac OS X Server requires detailed control over DHCP. DHCP is configured separately in Server Admin. The IP addresses given to VPN clients cannot overlap with addresses given to local DHCP clients. To learn more about DHCP, see Chapter 2, "DHCP Service," on page 23.

Enabling VPN also requires IP Firewall configuration. The firewall must be able to pass network traffic from external IP addresses through the firewall to the LAN. This can be as open or restricted as you deem necessary. For example, if the VPN clients are coming from a large range of IP addresses (you have many users each connecting from any number of ISPs) you may need to open the "any" firewall address group to VPN connections. If you want to narrow access to a small range of IP addresses, including static ones, you can create an address group that reflects that smaller range, and only enable VPN traffic originating from that list.

Managing VPN Service

This section describes tasks associated with managing VPN service. It includes starting, stopping, and configuring the service.

Starting or Stopping VPN Service

You use Server Admin to start and stop VPN service.

To start or stop VPN service:

- 1 In Server Admin, choose the VPN Service from the Computers & Services list.
- 2 Make sure at least one of the transport protocols is checked and configured.
- 3 Click Start Service or Stop Service.

When the service is turned on, the Stop Service button is available.

Enabling and Configuring L2TP Transport Protocol

Use Server Admin to designate L2TP as the transport protocol. By enabling this protocol, you must also configure the connection settings. You must designate an IPSec shared secret (if you don't use a signed Security Certificate), the IP address allocation range to be given to your clients, and group to be allowed VPN privileges (if desired). If both L2TP and PPTP are used, each protocol should have a separate, nonoverlapping address range.

To enable L2TP:

- 1 In Server Admin, choose the VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the L2TP tab.
- 4 Select "Enable L2TP over IPSec."
- 5 Set the beginning IP address of the allocation range.
- 6 Set the ending IP address of the allocation range.
- 7 Choose a PPP Authentication type.

If your computer is bound to a Kerberos authentication server, choose Kerberos, otherwise choose MS-CHAPv2.

- 8 Enter the shared secret, or select the certificate to use.
- 9 Click Save.

Enabling and Configuring PPTP Transport Protocol

Use Server Admin to designate PPTP as the transport protocol. By enabling this protocol, you must also configure the connection settings. You should designate an encryption key length (40-bit in addition to 128-bit), the IP address allocation range to be given to your clients, and group to be allowed VPN privileges (if desired). If both L2TP and PPTP are used, each protocol should have a separate, nonoverlapping address range.

To enable PPTP:

- 1 In Server Admin, choose the VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the PPTP tab.
- 4 Select “Enable PPTP.”
- 5 If desired, select “Allow 40-bit encryption keys” to allow such keys to be used in addition to 128-bit keys.

Warning: Allowing 40-bit encryption keys are much less secure, but may be necessary for some VPN client applications.

- 6 Set the beginning and IP addresses of the allocation range.
- 7 Click Save.

Configuring Additional Network Settings for VPN Clients

When a user connects in to your server through VPN, that user is given an IP address from your allocated range. This range is not served by a DHCP server, so you’ll need to configure additional network settings. These setting include the network mask, DNS address, and search domains.

To configure addition network settings:

- 1 In Server Admin, choose the VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Client Information tab.
- 4 Enter the IP address of the DNS server.
- 5 Enter any search domains, as needed.
- 6 Click Save.

Configuring VPN Network Routing Definitions

Network routing definitions allow you to choose whether to route data from the VPN clients to some address group either through the VPN tunnel (“private”) or over the VPN user’s ISP connection (“public”). For example, you may want all the VPN’s client traffic that goes to the LAN IP address range to go through the secure tunnel to the LAN, but make all traffic to other addresses to be routed through the user’s normal, unsecured Internet connection. This helps you have a finer control over what goes through the VPN tunnel.

Important Notes About VPN Routing Definitions:

- All traffic is routed through the VPN connection by default if no routing definitions are added.
- If any routing definitions are added, the VPN connection is no longer set as the default route, and any traffic destined for addresses not specifically declared as a Private route, will not go over the VPN connection.
- All DNS lookups go over the VPN connection at present regardless of the routes that are set.
- The definitions are unordered; they apply only the description that most closely matches the packet being routed.

An Example

Let's say your LAN's IP addresses are all 17.x.x.x addresses. If you make no routing definitions, every VPN client's network traffic (web browser URL requests, LPR printer queue print jobs, file server browsing) will be routed from his or her computer through the VPN tunnel to the 17.x.x.x LAN.

Now you decide that you don't want to handle all the traffic to web sites or file servers that aren't located on your network. You can restrict what traffic gets sent to the 17.x.x.x network, and what will go through the client's normal Internet connection. To limit the traffic the tunnel handles, you'd enter a routing definition designating traffic to the 17.x.x.x network as Private which sends it across the tunnel. In the routing definition table you'd enter:

```
17.0.0.0 255.0.0.0 Private
```

Now all traffic to the LAN is sent over the VPN connection, and by default, all other addresses not in the definitions table are sent over the users unencrypted Internet connection.

Now you decide that there a few IP addresses in the 17.x.x.x range that you don't want accessed over the VPN connection. You want the traffic to go via the user's own Internet connection, and not pass through the tunnel. The addresses might be outside the firewall, and not accessible from the 17.x.x.x LAN. As an example, let's use the addresses in the range 17.100.100.x. You would enter an additional routing definition as such:

```
17.100.100.0 255.255.255.0 Public
```

Since the address definition is more specific than 17.x.x.x, this rule takes precedence over the broader, more general rule, and traffic heading to any address in the 17.100.100.x range is sent over the VPN user's own Internet connection.

In summary, *if you add routes*, any routes you specify as Private will go over the VPN connection and any declared as Public will not go over the VPN connection. All others not specified will also *not* go over the VPN connection.

To set routing definitions:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Client Information tab.
- 4 Click the Add button below the routing definition list.
- 5 Enter a destination address range of the packets to be routed by specifying:
 - a A base address (for example, 192.168.0.0)
 - b A network mask (for example, 255.255.0.0)
- 6 Select the routing destination from the pop-up menu.
 - a Private means to route it through the VPN tunnel.
 - b Public means to use the normal interface with no tunnel.
- 7 Click OK.
- 8 Click Save.

Limiting VPN Access to Certain Users or Groups

By default, all users on the server or in the master directory have access to VPN, once enabled. You may want to limit VPN access to only certain users for security or ease of administration. You can limit access to VPN by using Mac OS X Server's Access Control Lists feature.

Access Control Lists (ACLs) are a method of designating service access to certain users or groups on an individual basis. For example, you may use an ACL to allow only one user access to a file server or shell login, without allowing any user on the server to access it.

To limit VPN access by login using ACLs:

- 1 In Server Admin, select the server which has VPN service running and the user or group that will have VPN access.
- 2 Click Access.
- 3 Deselect "Use same access for all services."
- 4 Select "Allow only users and group below."
- 5 Click the Add (+) button to reveal a Users and Groups drawer.
- 6 Drag the desired user or group to the access list.
- 7 Click Save.

Limiting VPN Access to Certain Incoming IP Addresses

By default, the IP Firewall blocks all incoming VPN connections. You may want to limit VPN access to only certain IP addresses for security or ease of administration. You can limit access to VPN by using configuring Mac OS X Server's IP Firewall feature.

To limit VPN access by IP address:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Advanced Rules tab.
- 4 Click the Add (+) button.
- 5 Select "allow" access in the Action pop-up menu.
- 6 Choose a from the Protocol pop-up menu.
If you have chosen L2TP for your VPN access, choose UDP.
If you have chosen PPTP for your VPN access, choose TCP.
- 7 Choose a VPN service type from the pop-up menu, L2TP or PPTP.
The appropriate destination port will automatically be added.
- 8 If desired, choose to log packets that match this filter rule.
- 9 Enter the source IP address range (with CIDR notation) you want to give access to the VPN, and leave the pop-up menu as Other.
These will be the IP addresses that will be able to connect to VPN service.
- 10 Choose the address group from the pop-up menu that has the VPN server for the destination of filtered traffic.
If you don't want to use an existing address group, enter the destination IP address range (with CIDR notation).
- 11 Choose the "In" network interface to apply this rule to.
"In" refers to the designated WAN interface.
- 12 Click OK.
- 13 Click Save to apply the filter immediately.

Additional Configuration Instructions

The following section has instructions for a few additional, optional scenarios. They require integration with an existing directory service system or third-party authentication services.

Enabling VPN-PPTP Access for Users in an LDAP Domain

In Mac OS X 10.4, you can use a command-line tool to enable PPTP-VPN connections for users who are in an LDAP domain.

This resolves a situation in which users can establish a VPN connection via PPTP to a Mac OS X Server that, once established, is not used by any network traffic. This affects Mac OS X Server 10.3 and 10.4.

- 1 Run the tool `/usr/sbin/vpnaddkeyagentuser` as root with the LDAP node (directory in which users are present) name as the argument.

For example, if the server that's running the VPN Service is also the LDAP Master, you would enter this command in Terminal:

```
sudo /usr/sbin/vpnaddkeyagentuser /LDAPv3/127.0.0.1
```

If the server that's running the VPN Service is not an LDAP Master, and the LDAP directory is on a different computer, use the IP address of the LDAP server in the command. For example, if the LDAP server is at 17.221.67.87, enter this command in Terminal:

```
sudo /usr/sbin/vpnaddkeyagentuser /LDAPv3/17.221.67.87
```

- 2 The tool will prompt for username and password.
 - a If the VPN Server is the LDAP master, type in the administrator name and password of the server.
 - b If the LDAP directory is on a different server, type in the administrator name and password of the server that hosts the LDAP directory (or the administrator name and password that is used to add users to the LDAP directory in Workgroup Manager). The tool will add a user to the LDAP directory and set up additional configuration elements in the VPN Server so that it can support PPTP.
- 3 Configure PPTP in the VPN Service Settings panel of the Server Admin.
- 4 Start VPN Service.

Offering SecurID Authentication With VPN Server

RSA Security provides strong authentication through their product offering. They use hardware and software tokens to verify user identity. SecurID authentication is available for both L2TP and PPTP transports. For details and product offerings, see: www.rsasecurity.com

Mac OS X Server VPN service can offer SecurID authentication, but it cannot be set up from within the Server Admin application. You can use Server Admin to configure standard VPN services, but Server Admin does not have an interface for choosing your authentication method. If you need to designate an authentication scheme other than the default (RSA Security's SecurID, for example), you will need to change the VPN configuration manually.

Setting up for SecurID

- 1 To configure RSA Security's SecurID authentication, you must first copy the `sdconf.rec` file from your SecurID server to a new directory on your Mac OS X Server named `/var/ace`.

There are several ways you could do this. These steps illustrate one method:

- a At your server, open the Terminal (`/Applications/Utilities/`).
 - b Type `sudo mkdir /var/ace` and press Return.
 - c Enter your administrator password, and press Return.
 - d Click the Finder icon in the Dock.
 - e From the Go menu, choose Go to Folder.
 - f Type: `/var/ace`
 - g Click Go.
 - h Copy the `sdconf.rec` file from your SecurID server into the "ace" folder.
 - i You will see a dialog indicating that the "ace" folder cannot be modified. Click the Authenticate button to allow the copy.
- 2 Second, you configure the VPN service on your Mac OS X Server to enable EAP-SecurID authentication for the protocols you want to use it with. To use it with PPTP, execute these two commands in Terminal (they are each only one line):

```
# sudo serveradmin settings
  vpn:Servers:com.apple.ppp.pptp:PPP:AuthenticatorEAPPlugins:_array_index
: 0 = "EAP-RSA"
# sudo serveradmin settings
  vpn:Servers:com.apple.ppp.pptp:PPP:AuthenticatorProtocol:_array_index:0
= "EAP"
```

To use it with L2TP, execute these two commands in Terminal:

```
# sudo serveradmin settings
  vpn:Servers:com.apple.ppp.l2tp:PPP:AuthenticatorEAPPlugins:_array_index
: 0 = "EAP-RSA"
# sudo serveradmin settings
  vpn:Servers:com.apple.ppp.l2tp:PPP:AuthenticatorProtocol:_array_index:0
= "EAP"
```

This is all that is required to configure SecurID. The remainder of Mac OS X Server VPN service configuration may be done using the Server Admin application.

Monitoring VPN Service

This section describes tasks associated with monitoring a functioning VPN service. It includes accessing status reports, setting logging options, viewing logs, and monitoring connections.

Viewing a VPN Status Overview

The VPN Overview gives you a quick status report on your enabled VPN services. It tells you how many L2TP and PPTP clients you have connected, which authentication method is selected, and when the service was started.

To view the overview:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click the Overview button.

Setting the Log Detail Level for VPN Service

You can choose the level of detail you want to log for VPN service.

- Nonverbose will indicate conditions for which you need to take immediate action (for example, if the VPN service can't start up).
- Verbose will record all activity by the VPN service, including routine functions.

Nonverbose login is enabled by default.

To set VPN log detail:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.
- 4 Select Verbose to enable verbose logging, if desired.
- 5 Click Save.

Viewing the VPN Log

You'll need to monitor VPN logs to ensure smooth operation of your Virtual Private Network. The VPN logs can help you troubleshoot problems. The log view is the `/var/log/ppp/vpnd.log` file. You can further filter the rules with the text filter box.

To view the log:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click Logs.

Viewing VPN Client Connections

You can monitor VPN client connections to ensure secure access to the Virtual Private Network. The client connection screen allows you to see the user connected, the IP address that user is connection from, the IP address assigned by your network, and the type and duration of connection.

You can sort the list by clicking on the column headers.

To view client connections:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click Connections.

Common Network Administration Tasks That Use VPN

The following sections illustrate some common network administration tasks that use VPN service.

Linking a Computer at Home With a Remote Network

VPN allows you to link a computer to a remote network, and access it as if it were physically connected to the LAN. This example uses the following information:

- *The user can authenticate with a name and password.*
- *Desired VPN type:* L2TP
- *Shared Secret:* prDwkj49fd!254
- *Internet or Public IP address of the VPN gateway:* gateway.example.com
- *Private Network IP address range and netmask:* 192.168.0.0–192.168.0.255 (also expressed as 192.168.0.0/24 or 192.168.0.0:255.255.255.0)
- *DHCP starting and ending addresses:* 192.168.0.3–192.168.0.127
- *Private Network's DNS IP address:* 192.168.0.2

The result of this configuration is a VPN client, able to connect to a remote LAN via L2TP, with full access to the LAN.

Step 1: Configure VPN

- 1 In Server Admin, choose the VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select L2TP.
- 5 Enter the shared secret (prDwkj49fd!254).

The shared secret is a common password that authenticates members of the cluster. IPSec uses the shared secret as a preshared key to establish secure tunnels between the cluster nodes.

- 6 Set the beginning IP address of the VPN allocation range.
It can't overlap the DHCP allocation range, so enter 192.168.0.128
- 7 Set the ending IP address of the VPN allocation range.
It can't overlap the DHCP allocation range, so enter 192.168.0.255
- 8 Leave the group blank, so all workgroups will have access to VPN login.
- 9 Click Save.
- 10 Select the Client Information tab.
- 11 Enter the IP address of the internal LAN DNS server (192.168.0.2)
- 12 Leave the routing definitions empty.
All traffic from the client will be through the VPN tunnel.
- 13 Click Save.
- 14 Start the VPN service.

Step 2: Configure the Firewall

- 1 Create an address group for the VPN allocation range.
See "Creating an Address Group" on page 67 for more information.
- 2 Open the firewall to external VPN connections by enabling L2TP connections in the "any" address group.
See "Opening the Firewall for Standard Services" on page 68 for more information.
- 3 Configure the firewall for the VPN address group, allowing or denying ports and services as desired.
- 4 Save your changes and start or restart the firewall, as needed.

Step 3: Configure the Client

This client example will be a Mac OS X client using Internet Connect.

- 1 Open Internet Connect.
- 2 Choose File > New VPN Connection.
- 3 Select L2TP over IPSec.
- 4 Select Edit Configuration from the Configuration pop-up menu.
- 5 Enter the following configuration information:
 - a Server Name: gateway.example.com
 - b Account name: <the user's short name>
 - c Authentication: Use Password <user's password>
 - d Shared Secret: prDwkj49fd!254
- 6 Click OK.

The user is now ready to connect.

Accessing a Single Computing Asset Behind a Remote Network Firewall

Accessing a single computing asset behind a firewall differs from allowing a client to become a node on the remote network. In the previous example, the VPN user's computer becomes a full participant in the Remote LAN. In this new scenario, the asset to be accessed is a single file server, with the VPN user's computer having no other contact with the remote LAN. This scenario assumes all the information in the section "Linking a Computer at Home With a Remote Network" on page 110, and adds:

- *File Server IP address:* 192.168.0.15
- *File Server Type:* Apple File Sharing

For this scenario, follow all the instructions in "Linking a Computer at Home With a Remote Network" on page 110, with these exceptions:

- In Step 1, part 12, don't leave the routing definitions empty.

Create a Private route with the IP number of the file server (192.168.0.15 255.255.255.255)

- In Step 2, part 3, configure the firewall to only accept Apple File Sharing Protocol connections and DNS from the VPN address group.

VPN users now logged in to through the VPN gateway will have access to the file server, and no other network traffic will go through the encrypted gateway.

Linking Two or More Remote Network Sites

VPN allows you to link not just a computer to a main network, but another network as well. This allows the two networks to interact as if they were physically connected together. Each site needs its own connection to the Internet, but the private data is sent encrypted between the two sites. A common use for this feature is to link your satellite offices to your organization's main office LAN.

About The Site-To-Site VPN Administration Tool

Linking multiple remote LAN sites to a main LAN requires the use of a command-line utility installed on Mac OS X Server called `s2svpnadmin` ("site-to-site VPN admin"). Using `s2svpnadmin` requires use of, and facility with, the Terminal as well as administrative access to root privileges through `sudo`. To find out more about `s2svpnadmin`, see its man page at:

```
man s2svpnadmin
```

Linking multiple remote LAN sites to a main LAN may also require the creation of a security certificate. The tool `s2svpnadmin` can create links using either shared-secret authentication (both sites have a password in their configuration files), or certificate authentication. If you want to use certificate authentication, you must create the certificate before running `s2svpnadmin`.

The site-to-site VPN connections can be only made using L2TP/IPSec VPN connections. You will not be able to link two sites using PPTP and these instructions.

This example uses the following information:

- *Desired VPN type:* L2TP
- *Authentication via shared secret.*
- *Shared Secret:* prDwkj49fd!254
- *Internet or Public IP address of the VPN main LAN gateway ("Site 1"):* A.B.C.D
- *Internet or Public IP address of the VPN remote LAN gateway ("Site 2"):* W.X.Y.Z
- *Private IP address of Site 1:* 192.168.0.1
- *Private IP address of Site 2:* 192.168.20.1
- *Private Network IP address range and netmask for Site 1:* 192.168.0.0–192.168.0.255 (also expressed as 192.168.0.0/24 or 192.168.0.0:255.255.0.0)
- *Private Network IP address range and netmask for Site 2:* 192.168.20.0–192.168.20.255 (also expressed as 192.168.20.0/16 or 192.168.0.0:255.255.0.0)
- *Organization's DNS IP address:* 192.168.0.2

The result of this configuration is an auxiliary, remote LAN, connected to a main LAN via L2TP.

Step 1: Run s2svpnadmin On Both Site's Gateways

- 1 In the Terminal, start s2svpnadmin by typing:

```
sudo s2svpnadmin
```
- 2 Enter the appropriate number to "Configure a new site-to-site server."
- 3 Enter an identifying configuration name (no spaces allowed).
For this example, you could enter "site_1" on Site 1's gateway, and so on.
- 4 Enter the gateway's public IP address.
For this example, enter A.B.C.D on Site 1's gateway, and W.X.Y.Z on Site 2's gateway.
- 5 Enter the other site's public IP address.
For this example, enter W.X.Y.Z on Site 1's gateway, and A.B.C.D on Site 2's gateway.
- 6 Enter "s" for shared secret authentication, and enter the shared secret: ("prDwkj49fd!254").
If you are using certificate authentication, enter "c" and choose the installed certificate to use.
- 7 Enter at least one addressing policy for the configuration.
- 8 Enter a local subnet address (for example, 192.168.0.0 for Site 1, 192.168.20.0 for Site 2)
- 9 Enter the prefix bits for the address range in CIDR notation.
In this example, the CIDR notation for the subnet range is 192.168.2.0/16 for Site 1, so you would enter "16."
- 10 Enter a remote subnet address (for example, 192.168.20.0 for Site 1, 192.168.0.0 for Site 2)

- 11 Enter the prefix bits for the address range in CIDR notation.
In this example, the CIDR notation for the subnet range is 192.168.2.0/16 for Site 1, so you would enter "16."
- 12 If you want to make more policies, indicate it now, otherwise press Return.
If you had more sites to connect, or a more complex address setup (linking only parts of your main LAN and the remote LAN) you would make more policies for this configuration now.
You'd repeat the previous policy steps for the new policies.
- 13 Enable the site configuration by pressing "y".
You can double check your settings by choosing to show the configuration details of the server, and entering the configuration name (in this example, "site_1").
- 14 Exit s2svpnadmin.

Step 2: Configure the Firewall On Both Site's Gateways

- 1 Create an address group with only the LAN gateway's public IP address.
In this example, use A.B.C.D/32 for Site 1, and W.X.Y.Z/32 for Site 2.
See "Creating an Address Group" on page 67 for more information.
- 2 Open the firewall to external VPN connections by enabling L2TP connections in the "any" address group.
See "Opening the Firewall for Standard Services" on page 68 for more information.
- 3 Create the following Advanced IP filter rules on both site's gateway's:

Rule 1

Action: Allow
Protocol: UDP
Source Address: A.B.C.D
Destination Address: W.X.Y.Z
Interface: Other, enter "isakmp"

Rule 2

Action: Allow
Protocol: UDP
Source Address: W.X.Y.Z
Destination Address: A.B.C.D
Interface: Other, enter "isakmp"

Rule 3

Action: Allow

Protocol: Other, enter "esp"

Source Address: A.B.C.D

Destination Address: W.X.Y.Z

Rule 4

Action: Allow

Protocol: Other, enter "esp"

Source Address: W.X.Y.Z

Destination Address: A.B.C.D

Rule 5

Action: Allow

Protocol: Other, enter "ipencap"

Source Address: A.B.C.D

Destination Address: W.X.Y.Z

Rule 6

Action: Allow

Protocol: Other, enter "ipencap"

Source Address: W.X.Y.Z

Destination Address: A.B.C.D

See "Creating an Advanced IP Firewall Rule" on page 70 for more information about creating advanced rules.

- 4 These rules will allow the encrypted traffic to be passed to both hosts.
- 5 Save your changes and start or restart the firewall, as needed.

Step 3: Start VPN Service On Both Site's Gateways

- 1 For both VPN gateways, in Server Admin, choose the VPN Service from the Computers & Services list.

If you've used s2svpnadmin correctly, the Start Service button should be enabled and ready to use.

- 2 Click Start Service.

You should now be able to access a computer on the Remote LAN from the local LAN. You could use `ping` or some means to verify the link.

Where to Find More Information

For more information about L2TP/IPSec:

The Internet Engineering Task Force (IETF) is working on formal standards for L2TP/IPsec user authentication. For more information, see the website:
www.ietf.org/ids.by.wg/ipsec.html

Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at the website www.ietf.org/rfc.html.

- For L2TP description, see RFC 2661.
- For PPTP description, see RFC 2637.
- For Kerberos version 5, see RFC 1510.

Network Time Protocol (NTP) is a network protocol used to synchronize the clocks of computers on your network to a time reference clock. NTP is used to ensure that all the computers on a network are reporting the same time.

If an isolated network, or even a single computer, is running on wrong time, services that use time and date stamps (like mail service, or web service with timed cookies) will send wrong time and date stamps and be out of synchronization with other computers across the Internet. For example, an email message could arrive minutes or years before it was sent (according to the time stamp), and a reply to that message could come before the original was sent.

How NTP Works

NTP uses Universal Time Coordinated (UTC) as its reference time. UTC is based on an atomic resonance, and clocks that run according to UTC are often called “atomic clocks.”

Internet-wide, authoritative NTP servers (called *Stratum 1* servers) keep track of the current UTC time. Other subordinate servers (called *Stratum 2 and 3* servers) query the Stratum 1 servers on a regular basis and estimate the time taken across the network to send and receive the query. They then factor this estimate with the query result to set the Stratum 2 or 3 servers’ own time. The estimates are accurate to the nanosecond.

Your local network can then query the Stratum 3 servers for the time. Then it repeats the process. An NTP client computer on your network then takes the UTC time reference and converts it, through its own time zone setting to local time, and sets its internal clock accordingly.

Using NTP on Your Network

Mac OS X Server can act not only as an NTP client, receiving authoritative time from an Internet time server, but also as an authoritative time server for a network. Your local clients can query your server to set their clocks. It's advised that if you set your server to answer time queries, you should also set it to query an authoritative server on the Internet.

Setting Up NTP Service

If you choose to run NTP service on your network, make sure your designated server can access a higher-authority time server. Apple provides a Stratum 2 time server for customer use at time.apple.com.

Additionally, you'll need to make sure your firewall allows NTP queries out to an authoritative time server on UDP port 123, and incoming queries from local clients on the same port. See Chapter 4, "IP Firewall Service," on page 59 for more information on configuring your firewall.

To set up NTP service:

- 1 Open Server Admin.
- 2 Make sure your server is configured to "Set Date & Time automatically."

This setting is in the Date & Time pane of the Server Admin Settings pane for the server.

- 3 Select the server you want to act as a time server.
- 4 Click Settings.
- 5 Select the General tab.
- 6 Select Enable NTP.
- 7 Click Save.

Configuring NTP on Clients

If you have set up a local time server, you can configure your clients to query your time server for getting the network date and time. By default, clients can query Apple's time server. These instructions allow you to set your clients to query your time server.

To configure NTP on clients:

- 1 Open System Preferences.
- 2 Click Date & Time.
- 3 Select "Set Date & Time automatically."
- 4 Select and delete the text in the field rather than use the pop-up menu.
- 5 Enter the host name of your time server.
Your host name can be either a domain name (like time.example.com) or an IP address.
- 6 Quit System Preferences.

Where to Find More Information

The NTP working group, documentation, and an F.A.Q. for NTP can be found at the website:

www.ntp.org

Listings of publicly accessible NTP servers and their use policies can be found at the website:

www.eecis.udel.edu/~mills/ntp/servers.html

Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at the website:

www.ietf.org/rfc.html

The official specification of NTP version 3 is RFC 1305.

Understanding 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 that are 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.

Setting Up Client Membership to a VLAN

You use the VLAN area of the Network pane of System Preferences to set up and manage VLANs. It's important to make sure ports with non-VLAN devices (non-802.1Q-compliant) are configured to transmit untagged frames. Many Ethernet cards are not 802.1Q-compliant. If they receive a tagged frame, they will not understand the VLAN tag and will drop the frame.

Note: This part of the Network pane is visible only if your hardware, such as an Xserve G5 system, supports this feature.

To configure a VLAN

- 1 Log in to your server as an administrator.
- 2 Open the Network pane of System Preferences.
- 3 Choose Network Port Configurations from the Show pop-up menu.
- 4 Click the VLAN button.
- 5 Select the Ethernet port you want to use for the VLAN.
- 6 Click Create VLAN.

- 7 Type a name for the VLAN, enter a tag (a number between 1 and 4094) in the Tag field, and click OK.

This VLAN tag designates the VLAN ID (VID). Each logical network has a unique VID. Interfaces that are configured with the same VID are on the same virtual network.

- 8 To use the VLAN, select it in the network port configurations list and click Apply Now.

Where to find more information

For more information about VLANs on the Internet:

www.ieee.org

VLAN standard is defined by IEEE.

Reference Document

The reference documents provides an overview of a protocol and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in a reference document helpful. If you're an experienced server administrator, you can find out technical details about a protocol in its reference document. It's available at:

standards.ieee.org/getieee802/download/802.1Q-1998.pdf

IPv6 is short for “Internet Protocol Version 6. IPv6 is the Internet’s next-generation protocol designed to replace the current Internet Protocol, IP Version 4 (IPv4, or just IP).

The current Internet Protocol is beginning to have problems coping with the growth and popularity of the Internet. IPv4’s main problems are:

- Limited IP addressing.
IPv4 addresses are 32 bits, meaning there can be only 4,300,000,000 network addresses.
- Increased routing and configuration burden.
The amount of network overhead, memory, and time to route IPv4 information is rapidly increasing with each new computer connected to the Internet.
- End-to-end communication is routinely circumvented.
This point is actually an outgrowth from the IPv4 addressing problem. As the number of computers increased and address shortages became more acute, another addressing and routing service was developed: Network Address Translation (NAT). NAT mediates and separates the two network end points. However, this frustrates a number of network services and is limiting.

IPv6 fixes some of these problems and helps others. It 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.

This chapter lists the IPv6 enabled services used by Mac OS X Server, gives guidelines for using the IPv6 addresses in those services, and explains IPv6 address types and notation.

IPv6 Enabled Services

The following services in Mac OS X Server support IPv6 in addressing:

- DNS (BIND)
- IP Firewall
- Mail (POP/IMAP/SMTP)
- SMB/CIFS
- Web (Apache 2)

Additionally, there are a number of command-line tools installed with Mac OS X Server that support IPv6 (for example, ping6, and traceroute6).

IPv6 Addresses in the Server Admin

The services above don't support IPv6 addresses in the user interface. They can be configured with command-line tools to add IPv6 addresses, but those same addresses will fail if entered into address fields in Server Admin.

IPv6 Addresses

IPv6 addresses are different than IPv4 addresses. In changing addresses, there are changes in address notation, reserved addresses, the address model, and address types.

Notation

While IPv4 addresses are 4 bytes long and expressed in decimals; IPv6 addresses are 16 bytes long and can be expressed a number of ways.

IPv6 addresses are generally written in the following form:

```
xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx
```

Pairs of IPv6 bytes are separated by a colon and each byte is represented as a pair of hexadecimal number, as in the following example:

```
E3C5:0000:0000:0000:4AC8:C0A8:6420
```

or

```
E3C5:0:0:0:4AC8:C0A8:6420
```

IPv6 addresses often contain many bytes with a zero value, so a shorthand notation is available. The shorthand notation removes the zero values from the text representation and puts the colons next to each other, as follows:

```
E3C5::4AC8:C0A8:6420
```

The final notation type includes IPv4 addresses. Because many IPv6 addresses are extensions of IPv4 addresses, the right-most 4 bytes of an IPv6 address (the right-most 2-byte pairs) can be rewritten in the IPv4 notation. This mixed notation (from the above example) could be expressed as:

E3C5:4AC8:192.168.100.32

IPv6 Reserved Addresses

IPv6 reserves two addresses that network nodes can't use for their own communication purposes:

0:0:0:0:0:0:0 (unspecified address, internal to the protocol)

0:0:0:0:0:0:0:1 (loopback address, just like 127.0.0.1 in IPv4)

IPv6 Addressing Model

IPv6 addresses are assigned to interfaces (for example, your Ethernet card), and not nodes (for example, your computer). A single interface can be assigned multiple IPv6 addresses. Also, a single IPv6 address can be assigned to several interfaces for load sharing. Finally, routers don't need an IPv6 address, eliminating the need to configure the routers for point-to-point unicasts. Additionally, IPv6 doesn't use IPv4 address classes.

IPv6 Address Types

IPv6 supports the following three IP address types:

- Unicast (one-to-one communication)
- Multicast (one-to-many communication)
- Anycast

Note that IPv6 does not support broadcast. Multicast is preferred for network broadcasts. Otherwise, unicast and multicast in IPv6 are the same as in IPv4. Multicast addresses in IPv6 start with "FF" (255).

Anycast is a variation of multicast. While multicast delivers messages to all nodes in the multicast group, anycast delivers messages to any one node in the multicast group.

Where to Find More Information

The working group for the Internet Protocol Version 6 website is:
www.ipv6.org

A group of IPv6 enthusiasts maintains a list of applications that support IPv6 at the website:
www.ipv6forum.com/navbar/links/v6apps.htm

Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at the website:
www.ietf.org/rfc.html

There are over 29 IPv6 related RFC documents. A list can be found at:
www.ipv6.org/specs.html

This glossary defines terms and spells out abbreviations you may encounter while working with online help or the Mac OS X Server Network Services Administration for Version 10.3 or Later manual. References to terms defined elsewhere in the glossary appear in *italics*.

access control A method of controlling which computers can access a network or network services.

access control list See **ACL**.

ACL Access Control List. A list maintained by a system that defines the rights of users and groups to access resources on the system.

address A number or other identifier that uniquely identifies a computer on a network, a block of data stored on a disk, or a location in a computer memory. See also **IP address**, **MAC address**.

bit A single piece of information, with a value of either 0 or 1.

bridge A computer networking device that connects two types of networking media, such as wireless and ethernet. It acts like a gateway by passing network traffic directly to the destination media without routing it or altering it in any way. Both sides of the network bridge need to have the same IP address subnet. It links small related network segments in a simple manner.

broadcast In a general networking context, the transmission of a message or data that any client on the network can read. Broadcast can be contrasted with unicast (sending a message to a specific computer) and multicast (sending a message to a select subset of computers). In QuickTime Streaming Server, the process of transmitting one copy of a stream over the whole network.

byte A basic unit of measure for data, equal to eight bits (or binary digits).

canonical name The “real” name of a server when you’ve given it a “nickname” or alias. For example, mail.apple.com might have a canonical name of MailSrv473.apple.com.

certificate Sometimes called an “identity certificate” or “public key certificate.” A file in a specific format (Mac OS X Server uses the x.509 format) that contains the public key half of a public-private keypair, the user’s identity information such as name and contact information, and the digital signature or either a *Certificate Authority* (CA) or the key user.

Certificate Authority An authority that issues and manages digital certificates in order to ensure secure transmission of data on a public network. See also *public key infrastructure and certificate*.

Challenge Handshake Authentication Protocol See **CHAP**.

CHAP Challenge Handshake Authentication Protocol. A common authentication protocol. See also **MS-CHAP**.

character A synonym for byte.

cleartext Data that hasn’t been encrypted.

command line The text you type at a shell prompt when using a command-line interface.

command-line interface A way of interfacing with the computer (for example, to run programs or modify file system permissions) by entering text commands at a shell prompt.

computer name The default name used for SLP and SMB/CIFS service registrations. The Network Browser in the Finder uses SLP to find computers advertising Personal File Sharing and Windows File Sharing. It can be set to bridge subnets depending on the network router settings. When you turn on Personal File Sharing, users see the computer name in the Connect To Server dialog in the Finder. Initially it is “<first created user>’s Computer” (for example, “John’s Computer”) but can be changed to anything. The computer name is used for browsing for network file servers, print queues, Bluetooth discovery, Apple Remote Desktop clients, and any other network resource that identifies computers by computer name rather than network address. The computer name is also the basis for the default **local hostname**.

cracker A malicious user who tries to gain unauthorized access to a computer system in order to disrupt computers and networks or steal information. Compare to hacker.

denial of service See **DoS attack**.

denial of service attack See **DoS attack**.

DHCP Dynamic Host Configuration Protocol. A protocol used to dynamically distribute IP addresses to client computers. Each time a client computer starts up, the protocol looks for a DHCP server and then requests an IP address from the DHCP server it finds. The DHCP server checks for an available IP address and sends it to the client computer along with a lease period—the length of time the client computer may use the address.

DHCP lease time See **lease period**.

directory services Services that provide system software and applications with uniform access to directory domains and other sources of information about users and resources.

DNS Domain Name System. A distributed database that maps IP addresses to domain names. A DNS server, also known as a name server, keeps a list of names and the IP addresses associated with each name.

DNS domain A unique name of a computer used in the Domain Name System to translate IP addresses and names. Also called a **domain name**.

DNS name A unique name of a computer used in the Domain Name System to translate IP addresses and names. Also called a **domain name**.

domain name See **DNS name**.

Domain Name System See **DNS**.

DoS attack Denial of service attack. An Internet attack that uses thousands of network pings to prevent the legitimate use of a server.

Dynamic Host Configuration Protocol See **DHCP**.

dynamic IP address An IP address that's assigned for a limited period of time or until the client computer no longer needs it.

EAP Extensible Authentication Protocol. An authentication protocol that supports multiple authentication methods.

encryption The process of obscuring data, making it unreadable without special knowledge. Usually done for secrecy and confidential communications. See also **decryption**.

Ethernet A common local area networking technology in which data is transmitted in units called packets using protocols such as TCP/IP.

Ethernet ID See **MAC address**.

filter A “screening” method used to control access to a server. A filter is made up of an IP address and a subnet mask, and sometimes a port number and access type. The IP address and the subnet mask together determine the range of IP addresses to which the filter applies.

firewall Software that protects the network applications running on your server. IP firewall service, which is part of Mac OS X Server software, scans incoming IP packets and rejects or accepts these packets based on a set of filters you create.

forward zone The DNS zone that holds no records of its own, but forwards DNS queries to another zone.

FTP File Transfer Protocol. A protocol that allows computers to transfer files over a network. FTP clients using any operating system that supports FTP can connect to a file server and download files, depending on their access privileges. Most Internet browsers and a number of freeware applications can be used to access an FTP server.

gateway A network node that interfaces one network to another. Often, it refers to a computer that links a private LAN to a public WAN, with or without Network Address Translation. A router is a special kind of gateway that links related network segments.

GB Gigabyte. 1,073,741,824 (2³⁰) bytes.

gigabyte See **GB**.

hacker An individual who enjoys programming, and explores ways to program new features and expand the capabilities of a computer system. See also **cracker**.

host name A unique name for a server, historically referred to as the UNIX hostname. The Mac OS X Server host name is used primarily for client access to NFS home directories. A server determines its host name by using the first name available from the following sources: the name specified in the /etc/hostconfig file (HOSTNAME=some-host-name); the name provided by the DHCP or BootP server for the primary IP address; the first name returned by a reverse DNS (address-to-name) query for the primary IP address; the local hostname; the name “localhost.”

HTTP Hypertext Transfer Protocol. The client/server protocol for the World Wide Web. The HTTP protocol provides a way for a web browser to access a web server and request hypermedia documents created using HTML.

Hypertext Transfer Protocol See **HTTP**.

IANA Internet Assigned Numbers Authority. An organization responsible for allocating IP addresses, assigning protocol parameters, and managing domain names.

ICMP Internet Control Message Protocol. A message control and error-reporting protocol used between host servers and gateways. For example, some Internet software applications use ICMP to send a packet on a round-trip between two hosts to determine round-trip times and discover problems on the network.

IEEE Institute of Electrical and Electronics Engineers, Inc. An organization dedicated to promoting standards in computing and electrical engineering.

IGMP Internet Group Management Protocol. An Internet protocol used by hosts and routers to send packets to lists of hosts that want to participate in a process known as multicasting. QuickTime Streaming Server (QTSS) uses multicast addressing, as does Service Location Protocol (SLP).

Internet Generally speaking, a set of interconnected computer networks communicating through a common protocol (TCP/IP). The Internet (note the capitalization) is the most extensive publicly accessible system of interconnected computer networks in the world.

Internet Assigned Numbers Authority See IANA.

Internet Control Message Protocol See ICMP.

Internet Group Management Protocol See IGMP.

Internet Message Access Protocol See IMAP.

Internet Protocol See IP.

Internet service provider See ISP.

IP Internet Protocol. Also known as IPv4. A method used with Transmission Control Protocol (TCP) to send data between computers over a local network or the Internet. IP delivers packets of data, while TCP keeps track of data packets.

IP address A unique numeric address that identifies a computer on the Internet.

IP subnet A portion of an IP network, which may be a physically independent network segment, that shares a network address with other portions of the network and is identified by a subnet number.

IPSec A security addition to IP. A protocol that provides data transmission security for L2TP VPN connections. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec nodes.

IPv4 See IP.

IPv6 Internet Protocol version 6. The next-generation communication protocol to replace IP (also known as IPv4). IPv6 allows a greater number of network addresses and can reduce routing loads across the Internet.

ISP Internet service provider. A business that sells Internet access and often provides web hosting for ecommerce applications as well as mail services.

KB Kilobyte. 1,024 (2¹⁰) bytes.

L2TP Layer Two Tunnelling Protocol. A network transport protocol used for VPN connections. It's essentially a combination of Cisco's L2F and PPTP. L2TP itself isn't an encryption protocol, so it uses IPSec for packet encryption.

LAN Local area network. A network maintained within a facility, as opposed to a WAN (wide area network) that links geographically separated facilities.

LDAP Lightweight Directory Access Protocol. A standard client-server protocol for accessing a directory domain.

lease period A limited period of time during which IP addresses are assigned. By using short leases, DHCP can reassign IP addresses on networks that have more computers than available IP addresses.

Lightweight Directory Access Protocol See **LDAP**.

list administrator A mailing list administrator. List administrators can add or remove subscribers from a mailing list and designate other list administrators. List administrators aren't necessarily local machine or domain administrators.

load balancing The process of distributing client computers' requests for network services across multiple servers to optimize performance.

local area network See **LAN**.

local domain A directory domain that can be accessed only by the computer on which it resides.

local hostname A name that designates a computer on a local subnet. It can be used without a global DNS system to resolve names to IP addresses. It consists of lowercase letters, numbers, or hyphens (except as the last characters), and ends with ".local" (e.g, bills-computer.local). Although the name is derived by default from the computer name, a user can specify this name in the Network pane of System Preferences. It can be changed easily, and can be used anywhere a DNS name or fully qualified domain name is used. It can only resolve on the same subnet as the computer using it.

MAC address Media access control address. A hardware address that uniquely identifies each node on a network. For AirPort devices, the MAC address is called the AirPort ID.

Mac OS X The latest version of the Apple operating system. Mac OS X combines the reliability of UNIX with the ease of use of Macintosh.

Mac OS X Server An industrial-strength server platform that supports Mac, Windows, UNIX, and Linux clients out of the box and provides a suite of scalable workgroup and network services plus advanced remote management tools.

mail exchange record See **MX record**.

Manual Unicast A method for transmitting a live stream to a single QuickTime Player client or to a computer running QTSS. An SDP file is usually created by the broadcaster application and then must be manually sent to the viewer or streaming server.

master zone The DNS zone records held by a primary DNS server. A master zone is replicated by zone transfers to slave zones on secondary DNS servers.

media access control See **MAC address**.

megabyte See **MB**.

Microsoft Challenge Handshake Authentication Protocol See **MS-CHAP**.

MS-CHAP Microsoft Challenge Handshake Authentication Protocol. The standard Windows authentication scheme for VPN. This authentication method encodes passwords when they are sent over the network and stores them in a scrambled form on the server. It offers good security during network transmission. MS-CHAP is a proprietary version of CHAP.

multicast In general, the simultaneous transmission of a message to a specific subset of computers on a network. See also **broadcast**, **unicast**. In QuickTime streaming, an efficient, one-to-many form of streaming. Users can join or leave a multicast but cannot otherwise interact with it.

multicast DNS A protocol developed by Apple for automatic discovery of computers, devices, and services on IP networks. This proposed Internet standard protocol is sometimes referred to as "ZeroConf." For more information, visit www.apple.com or www.zeroconf.org. To see how this protocol is used in Mac OS X Server, see **local hostname**.

multihoming The ability to support multiple network connections. When more than one connection is available, Mac OS X selects the best connection according to the order specified in Network preferences.

MX record Mail exchange record. An entry in a DNS table that specifies which computer manages mail for an Internet domain. When a mail server has mail to deliver to an Internet domain, the mail server requests the MX record for the domain. The server sends the mail to the computer specified in the MX record.

name server A server on a network that keeps a list of names and the IP addresses associated with each name. See also **DNS**, **WINS**.

NAT Network Address Translation. A method of connecting multiple computers to the Internet (or any other IP network) using one IP address. NAT converts the IP addresses you assign to computers on your private, internal network into one legitimate IP address for Internet communications.

NetInfo One of the Apple protocols for accessing a directory domain.

Network Address Translation See **NAT**.

network interface Your computer's hardware connection to a network. This includes (but isn't limited to) Ethernet connections, AirPort cards, and FireWire connections.

network interface card See **NIC**.

node A processing location. A node can be a computer or some other device, such as a printer. Each node has a unique network address. In Xsan, a node is any computer connected to a storage area network.

NTP Network time protocol. A network protocol used to synchronize the clocks of computers across a network to some time reference clock. NTP is used to ensure that all the computers on a network are reporting the same time.

Open Directory The Apple directory services architecture, which can access authoritative information about users and network resources from directory domains that use LDAP, NetInfo, or Active Directory protocols; BSD configuration files; and network services.

open relay A server that receives and automatically forwards mail to another server. Junk mail senders exploit open relay servers to avoid having their own mail servers blacklisted as sources of junk mail.

open source A term for the cooperative development of software by the Internet community. The basic principle is to involve as many people as possible in writing and debugging code by publishing the source code and encouraging the formation of a large community of developers who will submit modifications and enhancements.

packet A unit of data information consisting of header, information, error detection, and trailer records. QTSS uses TCP, UDP, and IP packets to communicate with streaming clients.

password An alphanumeric string used to authenticate the identity of a user or to authorize access to files or services.

password policy A set of rules that regulate the composition and validity of a user's password.

Password Server See **Open Directory Password Server**.

permissions Settings that define the kind of access users have to shared items in a file system. You can assign four types of permissions to a share point, folder, or file: read/write, read-only, write-only, and none (no access). See also **privileges**.

plaintext Text that hasn't been encrypted.

Point to Point Tunneling Protocol See **PPTP**.

pointer record See **PTR record**.

port A sort of virtual mail slot. A server uses port numbers to determine which application should receive data packets. Firewalls use port numbers to determine whether data packets are allowed to traverse a local network. "Port" usually refers to either a TCP or UDP port.

Post Office Protocol See **POP**.

PPTP Point to Point Tunneling Protocol. A network transport protocol used for VPN connections. It's the Windows standard VPN protocol and uses the user-provided password to produce an encryption key.

privileges The right to access restricted areas of a system or perform certain tasks (such as management tasks) in the system.

protocol A set of rules that determines how data is sent back and forth between two applications.

proxy server A server that sits between a client application, such as a web browser, and a real server. The proxy server intercepts all requests to the real server to see if it can fulfill the requests itself. If not, it forwards the request to the real server.

PTR record Pointer record. A DNS record type that translates IP (IPv4) addresses to domain names. Used in DNS reverse lookups.

QTSS QuickTime Streaming Server. A technology that lets you deliver media over the Internet in real time.

QuickTime Streaming Server See **QTSS**.

record type A specific category of records, such as users, computers, and mounts. For each record type, a directory domain may contain any number of records.

recursion The process of fully resolving domain names into IP addresses. A nonrecursive DNS query allows referrals to other DNS servers to resolve the address. In general, user applications depend on the DNS server to perform this function, but other DNS servers do not have to perform a recursive query.

relay In QuickTime Streaming Server, a relay receives an incoming stream and then forwards that stream to one or more streaming servers. Relays can reduce Internet bandwidth consumption and are useful for broadcasts with numerous viewers in different locations. In Internet mail terms, a relay is a mail SMTP server that sends incoming mail to another SMTP server, but not to its final destination.

scope A group of services. A scope can be a logical grouping of computers, such as all computers used by the production department, or a physical grouping, such as all computers located on the first floor. You can define a scope as part or all of your network.

search path See **search policy**.

search policy A list of directory domains searched by a Mac OS X computer when it needs configuration information; also the order in which domains are searched. Sometimes called a search path.

Secure Sockets Layer See **SSL**.

server A computer that provides services (such as file service, mail service, or web service) to other computers or network devices.

shared secret A value defined at each node of an L2TP VPN connection that serves as the encryption key seed to negotiate authentication and data transport connections.

shell A program that runs other programs. You can use a shell to interact with the computer by typing commands at a shell prompt. See also **command-line interface**.

shell prompt A character that appears at the beginning of a line in a command-line interface and indicates that you can enter a command.

SLP DA Service Location Protocol Directory Agent. A protocol that registers services available on a network and gives users easy access to them. When a service is added to the network, the service uses SLP to register itself on the network. SLP/DA uses a centralized repository for registered network services.

SMTP Simple Mail Transfer Protocol. A protocol used to send and transfer mail. Its ability to queue incoming messages is limited, so SMTP usually is used only to send mail, and POP or IMAP is used to receive mail.

spam Unsolicited email; junk mail.

SSL Secure Sockets Layer. An Internet protocol that allows you to send encrypted, authenticated information across the Internet. More recent versions of SSL are known as TLS (Transport Level Security).

SSL Secure Sockets Layer. An Internet protocol that allows you to send encrypted, authenticated information across the Internet. More recent versions of SSL are known as TLS (Transport Level Security).

static IP address An IP address that's assigned to a computer or device once and is never changed.

Stratum 1 An Internet-wide, authoritative network time protocol (NTP) server that keeps track of the current UTC time. Other strata are available (2, 3, and so forth); each takes its time from a lower-numbered stratum server.

subdomain Sometimes called the host name. Part of the domain name of a computer on the Internet. It does not include the domain or the top-level domain (TLD) designator (for example, .com, .net, .us, .uk). The domain name "www.example.com" consists of the subdomain "www," the domain "example," and the top level domain "com."

subnet A grouping on the same network of client computers that are organized by location (different floors of a building, for example) or by usage (all eighth-grade students, for example). The use of subnets simplifies administration. See also **IP subnet**.

subnet mask A number used in IP networking to specify which portion of an IP address is the network number.

TCP Transmission Control Protocol. A method used along with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. IP takes care of handling the actual delivery of the data, and TCP takes care of keeping track of the individual units of data (called packets) into which a message is divided for efficient routing through the Internet.

time server A network server to whose clock other computers on the network synchronize their own clocks, so that all computers report the same time. See also **NTP**.

time server A network server to whose clock other computers on the network synchronize their own clocks, so that all computers report the same time. See also **NTP**.

time-to-live See **TTL**.

Transmission Control Protocol See **TCP**.

TTL Time-to-live. The specified length of time that DNS information is stored in a cache. When a domain name-IP address pair has been cached longer than the TTL value, the entry is deleted from the name server's cache (but not from the primary DNS server).

TXT record Text record. A DNS record type that stores a text string for a response to a DNS query.

UCE Unsolicited commercial email. See **spam**.

UDP User Datagram Protocol. A communications method that uses the Internet Protocol (IP) to send a data unit (called a datagram) from one computer to another in a network. Network applications that have very small data units to exchange may use UDP rather than TCP.

unicast The transmission of data to a single recipient or client. If a movie is unicast to a user using RSTP, the user can move freely from point to point in an on-demand movie.

universal time coordinated See **UTC**.

User Datagram Protocol See **UDP**.

user name The long name for a user, sometimes referred to as the user's "real" name. See also **short name**.

UTC Universal time coordinated. A standard reference time. UTC is based on an atomic resonance, and clocks that run according to UTC are often called "atomic clocks."

Virtual Private Network See **VPN**.

VPN Virtual Private Network. A network that uses encryption and other technologies to provide secure communications over a public network, typically the Internet. VPNs are generally cheaper than real private networks using private lines but rely on having the same encryption system at both ends. The encryption may be performed by firewall software or by routers.

WAN Wide area network. A network maintained across geographically separated facilities, as opposed to a LAN (local area network) within a facility. Your WAN interface is usually the one connected to the Internet.

wildcard A range of possible values for any segment of an IP address.

Windows Internet Naming Service See **WINS**.

WINS Windows Internet Naming Service. A name resolution service used by Windows computers to match client names with IP addresses. A WINS server can be located on the local network or externally on the Internet.

WLAN A wireless local area network.

workgroup A set of users for whom you define preferences and privileges as a group. Any preferences you define for a group are stored in the group account.

zone transfer The method by which zone data is replicated among authoritative DNS servers. Slave DNS servers request zone transfers from their master servers to acquire their data.

A

AirPort Base Stations
DHCP service and 25

B

BIND 37, 38
load distribution 56

C

CIDR netmask notation 62, 64

D

DHCP servers 25
interactions 25
network location 25
DHCP service 23–36
AirPort Base Stations 25
changing subnets 27
deleting subnets 28
described 23
disabling subnets 28
DNS options 29
DNS Server for DHCP Clients 29
LDAP auto-configuration 25
LDAP options for subnets 29
logs 32
logs for 26
managing 26–31
more information 36
preparing for setup 23–25
setting up 26
starting and stopping 26
subnet IP addresses lease times, changing 28
subnet IP address lease times, changing 28
subnets 24
subnets, creating 27
subnet settings 27
uses for 23
viewing client lists 32
viewing leases, client list 32
WINS options for subnets 30, 31
Disabling 28

DNS service 37–58
described 37
load distribution 56
managing 41–42
more information 58
options for DHCP subnets 29
planning 38
preparing for setup 38
servers 38
setting up 38
setup overview 38–41
starting 41
stopping 41
strategies 38–41
uses for 37
with mail service 53
documentation 11
domain names
registering 38, 39
DoS (Denial of Service) attacks
preventing 80
Dynamic Host Configuration Protocol
See DHCP
dynamic IP addresses 24

F

filters
editing 71
examples 77–79
filters, IP
adding 65
described 62

I

IANA registration 38
Internet Gateway Multicast Protocol *See* IGMP
Internet Protocol Version 6 *See* IPv6
IP addresses
assigning 25
DHCP and 23
DHCP lease times, changing 28
dynamic 24
dynamic allocation 24

- IPv6 notation 124
 - leasing with DHCP 23
 - multiple 64
 - precedence in filters 64
 - ranges 64
 - reserved 25
 - static 24
 - IP Firewall
 - starting and stopping 31
 - IP Firewall service 59–61
 - about 59
 - adding filters 65
 - background 62
 - benefits 60
 - configuring 68–80
 - creating filters 69, 70
 - editing filters 71
 - example filters 77–79
 - filters 62–64
 - logs, setting up 74–75
 - managing 66–72
 - more information 86
 - multiple IP addresses 64
 - planning 65
 - port reference 82–86
 - preparing for setup 62–64
 - preventing Denial of Service (DoS) attacks 80
 - setting up 65–66
 - starting, stopping 66
 - uses for 60
 - viewing logs 74
 - IPv6
 - addressing 124–125
 - address notation 124
 - available services 124
 - in Server Admin 124
 - more information 126
- L**
- load distribution 56
 - logging items
 - DHCP activity 26
 - logs
 - DHCP 32
 - DNS service 49
 - IP Firewall service 74–76
- M**
- Mac OS X Server
 - ports used by 82–86
 - Mac OS X systems 82–86
 - mail
 - redirecting 53
 - Mail Exchange. *See* MX
 - mail exchangers 53
 - mail servers 53
 - mail service
 - using DNS service with 53
 - MX (Mail Exchange) records 40, 54
 - MX hosts 53
- N**
- name servers 38
 - NAT
 - about 87
 - configuring 89
 - monitoring 93
 - starting, stopping 89
 - status overview 93
 - NetBoot
 - viewing client lists 32
 - networks
 - private 56–57
 - TCP/IP networks 56–57
 - NTP
 - about 117
 - configuring clients 119
 - more information 119
 - setting up 118
 - time system 117
- P**
- ports
 - Mac OS X computers 82–86
 - TCP ports 82–83
 - UDP ports 85
- R**
- round robin 56
- S**
- Server 27, 32, 78, 79, 94, 95, 96, 97
 - server administration guides 11
 - servers
 - name servers 38
 - static IP addresses 24
 - Stratum time servers 117
 - subnet masks 62
 - subnets 24
 - creating 24, 27
- T**
- TCP/IP
 - private networks 56–57
 - TCP ports 82–84
 - time servers
 - Stratum 117
- U**
- UDP ports 85

Universal Time Coordinated (UTC) 117

V

VPN

client connections 110

logging 109

routing definitions 105

viewing logs 109

viewing status 109

