
Mac OS X Server
Network Services Administration
For Version 10.5 Leopard

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

019-0941/2007-09-01

Contents

About This Guide

This guide explains how to configure and administer Mac OS X Server network services.

Mac OS X Server version 10.5 includes several network services that help you manage and maintain your network.

What's New in Version 10.5

Mac OS X Server v10.5 offers the following major enhancements for network services:

- **New RADIUS feature:** Mac OS X Server v10.5 offers RADIUS for authorizing user access to AirPort Base Stations.
- **New services configuration assistants:** Mac OS X Server v10.5 offers a service configuration assistant for NAT and RADIUS.
- **Improved Bonjour:** Mac OS X Server v10.5 offers Bonjour administration.
- **Revised and improved firewall:** Mac OS X Server v10.5 uses an adaptive firewall that dynamically configures firewall rules and requires no configuration.

What's in This Guide

This guide includes the following chapters:

- Chapter 1, "Linking Your Network to the Internet," tells you how to use Gateway Setup Assistant to link your network to the Internet.
- Chapter 2, "Working with DHCP Service," tells you how to configure and use DHCP to assign IP addresses on your network.
- Chapter 3, "Working with DNS Service," tells you how to use Mac OS X Server as a domain name server.
- Chapter 4, "Working with Firewall Service," tells you how to maintain network security using a firewall.
- Chapter 5, "Working with NAT Service," tells you how to configure and use NAT to connect many computers to the Internet with only one public IP address.

- Chapter 6, "Working with VPN Service," tells you how to configure and use VPN to allow remote users to access your private LAN securely.
- Chapter 7, "Working with RADIUS Service," tells you how to configure and use RADIUS service to authorize Open Directory users and groups so they can access AirPort Base Stations on a network.
- Chapter 8, "Working with NTP Service," tells you how to enable your server as a time server.
- Chapter 9, "Supporting a VLAN," tells you about VLAN support for some server hardware configurations.
- Chapter 10, "Supporting IPv6," tells you about IPv6 and the services that support IPv6 addressing.

In addition, the Glossary provides brief definitions of the terms used in this guide.

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

Using This Guide

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

To get help for an advanced configuration of Leopard Server:

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

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

To see the most recent server help topics:

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

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

Mac OS X Server Administration Guides

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

www.apple.com/server/documentation.

This guide...	tells you how to:
Getting Started and *Mac OS X Server Worksheet*	Install Mac OS X Server and set it up for the first time.
Command-Line Administration	Install, set up, and manage Mac OS X Server using UNIX command-line tools and configuration files.
File Services Administration	Share selected server volumes or folders among server clients using the AFP, NFS, FTP, and SMB protocols.
iCal Service Administration	Set up and manage iCal shared calendar service.
iChat Service Administration	Set up and manage iChat instant messaging service.
Mac OS X Security Configuration	Make Mac OS X computers (clients) more secure, as required by enterprise and government customers.
Mac OS X Server Security Configuration	Make Mac OS X Server and the computer it's installed on more secure, as required by enterprise and government customers.
Mail Service Administration	Set up and manage IMAP, POP, and SMTP mail services on the server.
Network Services Administration	Set up, configure, and administer DHCP, DNS, VPN, NTP, IP firewall, NAT, and RADIUS services on the server.
Open Directory Administration	Set up and manage directory and authentication services, and configure clients to access directory services.
Podcast Producer Administration	Set up and manage Podcast Producer service to record, process, and distribute podcasts.

This guide...	tells you how to:
Print Service Administration	Host shared printers and manage their associated queues and print jobs.
QuickTime Streaming and Broadcasting Administration	Capture and encode QuickTime content. Set up and manage QuickTime streaming service to deliver media streams live or on demand.
Server Administration	Perform advanced installation and setup of server software, and manage options that apply to multiple services or to the server as a whole.
System Imaging and Software Update Administration	Use NetBoot, NetInstall, and Software Update to automate the management of operating system and other software used by client computers.
Upgrading and Migrating	Use data and service settings from an earlier version of Mac OS X Server or Windows NT.
User Management	Create and manage user accounts, groups, and computers. Set up managed preferences for Mac OS X clients.
Web Technologies Administration	Set up and manage web technologies, including web, blog, webmail, wiki, MySQL, PHP, Ruby on Rails, and WebDAV.
Xgrid Administration and High Performance Computing	Set up and manage computational clusters of Xserve systems and Mac computers.
Mac OS X Server Glossary	Learn about terms used for server and storage products.

Viewing PDF Guides on Screen

While reading the PDF version of a guide onscreen:

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

Printing PDF Guides

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

- Save ink or toner by not printing the cover page.
- Save color ink on a color printer by looking in the panes of the Print dialog for an option to print in grays or black and white.

- Reduce the bulk of the printed document and save paper by printing more than one page per sheet of paper. In the Print dialog, change Scale to 115% (155% for *Getting Started*). Then choose Layout from the untitled pop-up menu. If your printer supports two-sided (duplex) printing, select one of the Two-Sided options. Otherwise, choose 2 from the Pages per Sheet pop-up menu, and optionally choose Single Hairline from the Border menu. (If you're using Mac OS X v10.4 or earlier, the Scale setting is in the Page Setup dialog and the Layout settings are in the Print dialog.)

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

Getting Documentation Updates

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

- To view new onscreen help topics for a server application, make sure your server or administrator computer is connected to the Internet and click "Latest help topics" or "Staying current" in the main help page for the application.
- To download the latest guides in PDF format, go to the Mac OS X Server documentation website:

  www.apple.com/server/documentation

Getting Additional Information

For more information, consult these resources:

- *Read Me documents*—important updates and special information. Look for them on the server discs.
- *Mac OS X Server website* (www.apple.com/server/macosx)—gateway to extensive product and technology information.
- *Mac OS X Server Support website* (www.apple.com/support/macosxserver)—access to hundreds of articles from Apple's support organization.
- *Apple Training website* (www.apple.com/training)—instructor-led and self-paced courses for honing your server administration skills.
- *Apple Discussions website* (discussions.apple.com)—a way to share questions, knowledge, and advice with other administrators.
- *Apple Mailing Lists website* (www.lists.apple.com)—subscribe to mailing lists so you can communicate with other administrators using email.
- *OpenLDAP website* (www.openldap.org)—learn about the open source software that Open Directory uses to provide LDAP directory service.

- *MIT Kerberos website* (web.mit.edu/kerberos/www/)—get background information and specifications for the protocol that Open Directory uses to provide robust single sign-on authentication.
- *Berkeley DB website* (www.oracle.com/database/berkeley-db/)—investigate feature descriptions and technical documentation for the open source database that Open Directory uses to store LDAP directory data.
- *RFC3377, "Lightweight Directory Access Protocol (v3): Technical Specification"* (www.rfc-editor.org/rfc/rfc3377.txt)—lists a set of eight other Request for Comment (RFC) documents with overview information and detailed specifications for the LDAPv3 protocol.

Linking Your Network to the Internet

1

Use the Gateway Setup Assistant to guide you through the initial setup of your server to serve as a gateway between your private network and the Internet.

The Gateway Setup Assistant guides you through configuring your server to connect to the Internet. You make further changes to the service configuration using Server Admin. For network services, see the relevant section in this book for instructions.

About the Gateway Setup Assistant

The Gateway Setup Assistant helps you quickly and easily set up Mac OS X Server v10.5 to share your Internet connection with your local network. After you configure a few settings, the assistant can start sharing the server connection.

Depending on your configuration choices, the assistant performs the following when it sets up the server:

- Assigns the server a static IP address for each internal network interface.

  The address assigned is 192.168.x.1. The value used for x is determined by the network interface's order in the Network System Preference pane. For example, for the first interface on the list, x is 0; for the second interface, x is 1.
- Enables DHCP to allocate addresses on the internal network, removing existing DHCP subnets.
- Sets aside specific internal (192.168.x.x) addresses for DHCP use.

  Without VPN started, each interface can allocate addresses from 192.168.x.2 to 192.168.x.254.
- (Optional) Enables VPN to permit authorized external clients to connect to the local network.

  VPN L2TP is enabled, so you must enter a shared secret (a passphrase) for client connections to use.
- Sets aside specific internal addresses (192.168.x.x) 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.

- Enables the firewall to help secure the internal network.

  Address groups are added for each internal network interface, with all traffic permitted from the newly created DHCP address ranges to any destination address.
- Enables network address translation (NAT) on the internal network and adds a NAT divert rule to the IP firewall to direct network traffic to the correct computer. This also protects the internal network from unsolicited external connections.
- Enables DNS on the server, configured to cache lookups, to improve DNS response for internal clients.

Before configuring these settings, you can review the proposed changes before committing to them and overwriting existing settings.

You can make further changes to the service configuration using Server Admin. For network services, see the relevant section in this book for information.

If you run the Gateway Setup Assistant again, it overwrites manual settings you made.

Running the Gateway Setup Assistant

You run the Gateway Setup Assistant from the NAT service Overview pane in Server Admin.

To run the Gateway Setup Assistant:

1. Open Server Admin and connect to the server.
2. Click Settings, then click Services.
3. Select the NAT checkbox, then click Save.
4. Click the triangle to the left of the server.

   The list of services appears.
5. From the expanded Servers list, select NAT.
6. Click Overview.
7. Click Gateway Setup Assistant.
8. Follow the directions in the assistant, click Continue after each page, read the final configuration summary carefully, and make sure you approve of the settings before finalizing the configuration.

> ***WARNING:*** Although you can use Service Configuration Assistant to configure remote servers, you can accidentally cut off your administrator access to the remote server.

Connecting a Wired LAN to the Internet

You can use the Gateway Setup Assistant to connect a wired LAN to the Internet. Your LAN can consist of any number of computers connected to each other through Ethernet hubs and switches, but the LAN must have one point of contact with the Internet (the gateway).

Your gateway has one connection to the Internet and one connection to the LAN. All other computers access the Internet through your gateway. You can configure your Mac OS X server to be a gateway to the Internet, which requires that your server have two Ethernet ports (en0 and en1). Ethernet en0 should be connected to the Internet and en1 should be connected to your LAN.

After this process, computers on the LAN:

- Can get IP addresses and network settings that were configured using DHCP.
- Can access the Internet if the gateway is connected to the Internet.
- Can't be accessed by unauthorized network connections originating from the Internet.
- Can be accessed over the Internet by authorized VPN clients (if VPN is configured).
- Can benefit from DNS lookup caching in the gateway, which speeds DNS resolution.

To connect a wired LAN to the Internet:

1. Plug the connection to the Internet into the Ethernet 1 (en0) port.
2. Plug the connection to your LAN into the Ethernet 2 (en1) port.
3. Open Server Admin and connect to the server.
4. Click Settings, then click Services.
5. Select the NAT checkbox.
6. Click Save.
7. Click the triangle to the left of the server.

   The list of services appears.
8. From the expanded Servers list, select NAT.
9. Click Overview, then click Gateway Setup Assistant.
10. Click Continue.

    If your server has existing DHCP, DNS, NAT, and VPN configurations, you are prompted to overwrite those configurations. If you want to overwrite existing configurations, click Overwrite to continue.
11. From the Gateway WAN Interface pop-up menu, choose Ethernet 1 (en0) for you WAN interface, then click Continue.
12. From the list of network interfaces, select the Ethernet 2 checkbox for you LAN interface and click Continue.

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

If your server has more than one interface available (Ethernet port 2, Ethernet port 3, and so on), choose those you want to enable.

13 (Optional) If you want to make your gateway server a VPN entry point to your LAN, select the Enable VPN for this server checkbox.

If you enable VPN, you need a shared secret. A shared secret is a passphrase that users must provide to securely connect to the VPN gateway. It should be a very secure passphrase, not the password of a user or administrator on the gateway server.

To set a very secure passphrase, use Password Assistant in Account Preferences. For more information, see *Mac OS X Server Security Configuration*.

For more information, see Chapter 6, "Working with VPN Service."

14 Click Continue.

15 Inspect and confirm your setup.

16 Click Continue.

NAT and all dependent services will be configured and started.

17 Click Close.

Options

You can fine-tune the settings of this base configuration, but you perform additional configuration in Server Admin.

For example, you can use Server Admin to assign IP addresses to specific computers. To do this, add static address mappings in the DHCP service settings. For more information, see Chapter 2, "Working with DHCP Service."

You can also change firewall settings to permit connections from the Internet to the LAN. To do this, change the firewall settings, open up IP ports as needed, and configure port forwarding (by editing UNIX files from the command line) to designate which computer on the LAN is to accept incoming traffic.

Connecting a Wired LAN and Wireless Clients to the Internet

You can use the Gateway Setup Assistant to connect a wired LAN and wireless clients to the Internet. Your LAN can consist of any number of computers connected to each other through Ethernet hubs and switches, but the LAN must have 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 wired network. 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, computers on the LAN and those connected to the AirPort Base Station:

- Can get IP addresses and network settings configured using DHCP.
- Can access the Internet, if the gateway is connected to the Internet.
- Can't be accessed by unauthorized network connections originating from the wired connection to the Internet.
- Can be accessed over the Internet by authorized VPN clients (if VPN is configured).
- Can benefit from DNS lookup caching in the gateway, which speeds DNS resolution.

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

1. Plug the connection to the Internet into the Ethernet 1 (en0) port.
2. Plug the connection to your LAN into the Ethernet 2 (en1) port.
3. Connect the AirPort Base Station port (the WAN port, if there are two) to the wired network.
4. Using the AirPort Utility, configure the Base Station to connect using Ethernet and to get its address using DHCP.

   You can open it from the /Applications/Utilities/ folder.
5. Select your base station, and then choose Manual Setup from the Base Station menu.
6. Enter the base station password if necessary.
7. Click Internet in the toolbar, then click Internet Connection.
8. From the Connect Using pop-up menu choose Ethernet.
9. From the Configure IPv4 pop-up menu choose Using DHCP.
10. From the Connection Sharing pop-up menu choose Off (Bridge Mode).
11. To change Base Station settings, click Update.
12. Open Server Admin and connect to the server.
13. Click Settings, then click Services.
14. Select the NAT checkbox.
15. Click Save.
16. Click the triangle to the left of the server.

    The list of services appears.
17. From the expanded Servers list, select NAT.
18. Click Overview, then click Gateway Setup Assistant.
19. Click Continue.
20. For your WAN (Internet) interface, designate Ethernet 1.
21. For your LAN (sharing) interface, designate Ethernet 2.

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

If your server has more than one interface available (Ethernet port 2, Ethernet port 3, and so on), choose those you want to enable.

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

If you enable VPN, you need a shared secret. A shared secret is a passphrase that users must provide to securely connect to the VPN gateway. It should be a very secure passphrase, not a password of a user or administrator on the gateway server.

To set a very secure passphrase, use Password Assistant in Account Preferences. For more information, see *Mac OS X Server Security Configuration*.

For more information about VPN, see Chapter 6, "Working with VPN Service."

23 Inspect and confirm the changes.

Options

You can fine-tune the settings of this base configuration, but you perform additional configuration in Server Admin.

For example, you can use Server Admin to assign IP addresses to specific computers. To do this, add static address mappings in the DHCP section's Settings tab. For more information, see Chapter 2, "Working with DHCP Service."

You can also change firewall settings to permit connections from the Internet to the LAN. To do this, change the firewall settings, opening up IP ports as needed, and configure port forwarding in the NAT pane to designate which computer on the LAN is to accept incoming traffic.

Connecting a Wireless LAN to the Internet

Connecting wireless clients to the Internet through a Mac OS X Server gateway provides the following advantages over using AirPort Base Station built-in functions:

- Advanced firewall control
- DHCP allocation of static IP addresses
- DNS caching
- Incoming VPN connections to the LAN

If you do not need these advanced functions, use the AirPort Base Station to connect your wireless clients to the Internet without using 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 network traffic through the gateway.

All wireless clients must be able to connect to the AirPort Base Station's wireless network to be linked to the gateway.

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

- Can get IP addresses and network settings configured using DHCP.
- Can access the Internet if the gateway is connected to the Internet.
- Can't be accessed by unauthorized network connections originating from the wired connection to the Internet.
- Can be accessed over the Internet by authorized VPN clients (if VPN is configured).
- Can benefit from DNS lookup caching in the gateway, which speeds DNS resolution.

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

1. Plug the connection to the Internet into the Ethernet 1 (en0) port.
2. Connect the AirPort Base Station port (the WAN port, if there are two) to the Ethernet 2 (en1) port.
3. Using the AirPort Utility, configure the Base Station to connect using Ethernet and to get its address using DHCP.

   You can open it from the /Applications/Utilities/ folder.
4. Select your base station, and then choose Manual Setup from the Base Station menu.
5. Enter the base station password if necessary.
6. Click Internet in the toolbar, then click Internet Connection.
7. From the Connect Using pop-up menu choose Ethernet.
8. From the Configure IPv4 pop-up menu choose Using DHCP.
9. From the Connection Sharing pop-up menu choose Off (Bridge Mode).
10. To change Base Station settings, click Update.
11. Open Server Admin and connect to the server.
12. Click Settings, then click Services.
13. Select the NAT checkbox.
14. Click Save.
15. Click the triangle to the left of the server.

    The list of services appears.
16. From the expanded Servers list, select NAT.
17. Click Overview, then click Gateway Setup Assistant.
18. Click Continue.
19. For your WAN (Internet) interface, designate Built-In Ethernet 1.
20. For your LAN (sharing) interface, designate Built-In Ethernet 2.

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

If your server has more than one interface available (Ethernet port 2, Ethernet port 3, and so on), choose those you want to enable.

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

If you enable VPN, you need a shared secret. A shared secret is a passphrase that users must provide to securely connect to the VPN gateway. It should be a very secure passphrase, not a password of a user or administrator on the gateway server.

To set a very secure passphrase, use Password Assistant in Account Preferences. For more information, see *Mac OS X Server Security Configuration*.

For more information about VPN, see Chapter 6, "Working with VPN Service."

22 Inspect and confirm the changes.

Options

You can fine-tune the settings from this base configuration but you perform additional configuration in Server Admin.

For example, you can use Server Admin to assign IP addresses to specific computers. To do this, add static address mappings in the DHCP section's Settings tab. For more information, see Chapter 2, "Working with DHCP Service."

You can also change firewall settings to permit connections from the Internet to the LAN. To do this, change the firewall settings, opening up IP ports as needed, and configure port forwarding in the NAT pane to designate which computer on the LAN is to accept incoming traffic.

Working with DHCP Service

2

This chapter describes how to set up and manage DHCP service in Mac OS X Server.

If your organization has more clients than IP addresses, you can benefit from using Dynamic Host Configuration Protocol (DHCP) service. IP addresses are assigned as needed, and when they're not needed, they can be used by other clients. You can use a combination of static and dynamic IP addresses for your network.

DHCP service lets you administer and distribute IP addresses to 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 computer configured to use DHCP starts up, it looks for a DHCP server on your network. If it finds a DHCP server, the client computer then requests an IP address. The DHCP server checks for an available IP address and sends it to the computer with a *lease period* (the length of time the client computer can use the address) and configuration information.

For more information about static and dynamic allocation of IP addresses, see "Before Setting Up DHCP Service" on page 26.

Organizations can 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 computers without needing to configure each client.

You can use the DHCP module in Server Admin to:

- Configure and administer DHCP service
- Create and administer subnets
- Configure DNS, LDAP, and Windows Internet Naming Service (WINS) options for client computers
- View DHCP address leases

Setup Overview

Here is an overview of the basic steps for setting up DHCP service.

Note: If you used the Gateway Setup Assistant to configure ports on your server when you installed Mac OS X Server, some DHCP information is already configured. 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 31.

Step 1: Before you begin

For issues to keep in mind when you setup DHCP service, read "Before Setting Up DHCP Service" on page 26.

Step 2: Turn DHCP service on

Before configuring DHCP service, turn on DHCP. See "Turning DHCP Service On" on page 29.

Step 3: Create subnets

Use Server Admin 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 29.

Step 4: Configure DHCP log settings

You can log the activity and errors in your DHCP service to help you identify use patterns and problems with your server.

The 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. See "Configuring Log Settings" on page 30.

Step 5: Start DHCP service

After you configure DHCP, start the service to make it available. See "Starting DHCP Service" on page 30.

Before Setting Up DHCP Service

This section provides 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 a network that simplify administration. You can organize subnets any way that is useful to you. For example, you can create subnets for different groups in your organization or for different floors of a building.

After you group computers into subnets, you can configure options for all computers on a subnet at one time instead of setting options for individual computers.

Each subnet needs a way to connect to other subnets. A hardware device called a *router* typically connects subnets.

Assigning IP Addresses Dynamically

With dynamic address allocation, an IP address is assigned for a limited period of time (the lease time) or until the 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 IP addresses. Leases are renewed if the address isn't needed by another 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, "Working with VPN Service," on page 125.

Using Static IP Addresses

Static IP addresses are assigned to a computer or device once and then don't change. You can 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, can also benefit from static IP addresses.

Static IP addresses can be set up manually by entering the IP address on the computer (or other device) that is assigned the address, or by configuring DHCP to provide the same address to a specific computer or device on each request.

Manually configured static IP addresses avoid potential issues that some services can have with DHCP-assigned addresses, and they don't suffer from the delay that DHCP requires to assign an address.

DHCP-assigned addresses permit address configuration changes at the DHCP server rather than at each client.

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

You can set up DHCP to always serve the same address to the same computer. For more information, see "Assigning Static IP Addresses Using DHCP" on page 35.

Locating the DHCP Server

When a computer looks for a DHCP server, it broadcasts a message. If your DHCP server is on a different subnet from the computer, 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, place the DHCP server on the same subnet as your client.

Interacting with Other DHCP Servers

You might already have 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 might want your DHCP server to provide an LDAP server address for client autoconfiguration in managed environments.

Because AirPort Base Stations can't provide an LDAP server address, 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, your routers must be configured to forward client broadcasts and DHCP server responses as described previously.

To provide DHCP service with AirPort Base Stations. You must manually enter LDAP server addresses of computers. You can't use the client autoconfiguration feature.

Using Multiple DHCP Servers on a Network

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

Assigning Reserved IP Addresses

Some IP addresses can't be assigned, including addresses reserved for loopback and for broadcasting. Your ISP won't assign these addresses to you. If you try to configure DHCP to use these addresses, you're warned that the addresses are invalid and you must enter valid addresses.

Getting More Information About the DHCP Process

Mac OS X Server uses a daemon process named `bootpd` that is responsible for the DHCP Service's address allocation. For more information about `bootpd` and its advanced configuration options, see the `bootpd` man page.

Turning DHCP Service On

Before you can configure DHCP settings, you must turn on DHCP service in Server Admin.

To turn DHCP service on:

1. Open Server Admin and connect to the server.
2. Click Settings.
3. Click Services.
4. Select the DHCP checkbox.
5. Click Save.

Setting Up DHCP Service

Set up DHCP service by configuring the following items in Server Admin:

- **Subnet.** Create a pool of IP addresses that are shared by computers on your network.
- **Log Level.** Configure the DHCP event log level.

The following sections describe the tasks for configuring these settings. A final section tells you how to start DHCP service when you finish.

Creating Subnets in DHCP Service

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

To create a subnet:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Subnets.
5. Click the Add (+) button.
6. Enter a descriptive name for the new subnet.
7. Enter a starting and ending IP address for this subnet range.

   Addresses must be contiguous and they can't overlap with other subnet ranges.
8. Enter the subnet mask for the network address range.
9. From the pop-up menu, choose the network interface that will host DHCP service.
10. Enter the IP address of the router for this subnet.

If the server you're configuring 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 want to set DNS, LDAP, or WINS information for this subnet, enter these now.

For more information, see "Setting the DNS Server for a DHCP Subnet" on page 33, "Setting LDAP Options for a Subnet" on page 34, and "Setting WINS Options for a Subnet" on page 34.

13 Click Save.

14 To enable the subnet, select the Enable checkbox.

15 Click Save.

Configuring Log Settings

You can choose the level of detail you want for DHCP service logs:

- **Low (errors only)**: Indicates conditions where you must 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)**: Alerts you to conditions where data is inconsistent but the DHCP server can still operate. This level corresponds to default bootpd reporting.
- **High (all events)**: Records 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 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DHCP.

4 Click Settings.

5 From the Log Level pop-up menu, choose the logging option you want.

6 Click Save.

Starting DHCP Service

You start the DHCP service to provide IP addresses to your users. You must have at least one subnet created and enabled.

To start DHCP service:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DHCP.

4 Click the Start DHCP button (below the Servers list).

If the Firewall service is running, a warning appears asking you to verify that all ports used by DHCP are open. Click OK.

The service runs until you stop it. It restarts when your server is restarted.

From the Command Line

You can also start the DHCP service using the `serveradmin` command in Terminal. For more information, see the file services chapter of *Command-Line Administration*.

Managing DHCP Service

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

Stopping DHCP Service

When starting or stopping DHCP, you must have at least one subnet created and enabled.

To stop DHCP service:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DHCP.

4 Click the Stop DHCP button (below the Servers list).

5 Click Stop Now.

Changing Subnet Settings in DHCP Service

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

To change subnet settings:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DHCP.

4 Click Subnets.

5 Select a subnet.

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.

If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Deleting Subnets from DHCP Service

You can delete subnets and subnet IP address ranges so they are no longer distributed to computers.

To delete subnets or address ranges:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DHCP.

4 Click Subnets.

5 Select a subnet.

6 Click the Delete (–) button.

7 Click Save.

If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Disabling Subnets Temporarily

You can temporarily shut down a subnet without losing its settings. No IP addresses from the subnet's range are distributed on the selected interface to any computer until you reenable the subnet.

To disable a subnet:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DHCP.

4 Click Subnets.

5 Deselect Enable next to the subnet you want to disable.

6 Click Save.

If DHCP is running, you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Changing IP Address Lease Times for a Subnet

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

To change the lease time for a subnet:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Subnets.
5. Select a subnet.
6. From the Lease Time pop-up menu, choose a time scale (hours, days, weeks, or months).
7. In the Lease Time field, enter a number.
8. Click Save.

   If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Setting the DNS Server for a DHCP Subnet

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

To set DNS options for a subnet:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Subnets.
5. Select a subnet.
6. Click DNS.
7. Enter the primary and secondary name server IP addresses you want DHCP clients to use.
8. Enter the default domain of the subnet.
9. Click Save.

   If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Setting LDAP Options for a Subnet

You can use DHCP to automatically 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 are using this Mac OS X Server as an LDAP master, LDAP options are populated with the necessary configuration information. If your LDAP master server is another computer, you must know the domain name or IP address of the LDAP database that you want to use, and you must know the LDAP search base.

To set LDAP options for a subnet:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Subnets.
5. Select a subnet.
6. Click LDAP.
7. Enter the domain name or IP address of the LDAP server for this subnet.
8. Enter the search base for LDAP searches.
9. If you're using a nonstandard port, enter the LDAP port number.
10. If necessary, select LDAP over SSL.

    Use this option to secure LDAP communication.
11. Click Save.

    If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Setting WINS Options for a Subnet

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

You must know the domain name or IP address of the Windows Internet Naming Service/NetBIOS Name Server (WINS/NBNS) primary and secondary servers (usually the IP address of the DHCP server), and the NetBIOS over TCP/IP (NBT) node type.

The following are possible node types:

- **Hybrid (h-node):** Checks the WINS server and then broadcasts

- **Peer (p-node):** Checks the WINS server for name resolution
- **Broadcast (b-node):** Broadcasts for name resolution (most commonly used)
- **Mixed (m-node):** Broadcasts for name resolution and then checks the WINS server

The NetBIOS Datagram Distribution (NBDD) server works with NBNS to route datagrams to computers on a different subnet.

The NetBIOS Scope ID isolates NetBIOS communication on a network. The NetBIOS Scope ID is appended to the NetBIOS name of the computer. All computers that have the same NetBIOS Scope ID can communicate.

NBDD Server and the NetBIOS Scope ID are typically not used, but you might need to use them depending on your Windows clients' configuration and Windows network infrastructure.

To set WINS options for a subnet:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Subnets.
5. Select a subnet.
6. Click WINS.
7. Enter the domain name or IP address of the WINS/NBNS primary and secondary servers for this subnet.
8. Enter the domain name or IP address of the NBDD server for this subnet.
9. From the pop-up menu, choose the NBT node type.
10. Enter the NetBIOS Scope ID.
11. Click Save.

    If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Assigning Static IP Addresses Using DHCP

You can assign the same address to the same computers. This helps simplify configuration when using DHCP and lets you have some static servers or services.

To keep the same IP address for a computer, you must know the computer's Ethernet address (also known as the MAC address or hardware address). Each network interface has its own Ethernet address.

If a computer is connected to a wired network and a wireless network, it uses a different Ethernet address for each network connection.

To assign static IP addresses:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Static Maps.
5. Click Add Computer.
6. Enter the name of the computer.
7. In the Network Interfaces list, click the column to enter the following information:

   MAC Address of the computer that needs a static address.

   IP address you want to assign to the computer.
8. If your computer has other network interfaces that require static IP addresses, click the Add (+) button and enter the IP address you want to assign for each interface.
9. Click OK.
10. Click Save.

    If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Removing or Changing Static Address Maps

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

To change the static address map:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Static Maps.
5. Select a mapping to Edit or Remove.
6. Click the Edit button or the Remove button.

   If you are editing the mapping, make changes you want, then click OK.
7. Click Save.

   If DHCP is running you are prompted to restart DHCP for your change to take effect. Otherwise, your changes take effect the next time you start DHCP.

Monitoring DHCP Service

You can use the following methods to monitor and troubleshoot DHCP service:

- Monitor the computers that are using the service by viewing the client list.
- Monitor the log files generated by the service.
- Use service logs to troubleshoot network problems.

The following sections discuss these aspects of DHCP service.

Checking DHCP Service Status

The status overview shows the following summary of the DHCP service.

- Whether the service is running
- How many clients it has
- When the service was started
- How many IP addresses are statically assigned from your subnets
- The last time the client database was updated

To view DHCP service status:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Overview to view whether the service is running, when it started, the number of clients connected, and the when the last database update occurred.

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`. Use the Filter field to search for specific entries.

To view DHCP log entries:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Log.
5. To search for specific entries, use the Filter field (upper right corner).

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 (or the number of hours and minutes, if less than 24 hours)
- The DHCP client ID (usually the same as the hardware address)
- The computer name
- The hardware address

To view the DHCP client list:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Clients.

   To sort the list by different criteria, click a column heading.

Common Network Configurations That Use DHCP

The following section contains example DHCP configurations for different network uses. These include a workgroup configuration, a student lab configuration, and a coffee shop configuration.

When you set up a private network, you choose IP addresses from the blocks of IP addresses reserved by the Internet Assigned Numbers Authority (IANA) 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 Addresses Behind a NAT Gateway

You use DHCP to provide IP addresses to computers behind a Network Address Translation (NAT) gateway.

Although not strictly necessary (because NAT can be used with static IP addresses instead of DHCP), this enables easy configuration of computers.

For more information, see “Linking a LAN to the Internet Through One IP Address” on page 118.

Workgroup Configuration

Imagine you have a small workgroup with its own DHCP address group. You can have an IP-connected printer, a file server, and an Open Directory server (on or off the subnet) for user management purposes.

To use DHCP in this setting, you must already have:

- A working, configured firewall that permits LDAP and printer (IP printing) connections.

  For more information, see Chapter 4, "Working with Firewall Service."

- A working, configured Open Directory or LDAP server with users defined.

  For more information, see *Open Directory Administration* and *User Management*.

For this example, configuring DHCP involves 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 address of the printer. If the printer can be configured to accept an address using DHCP, don't worry about an overlap.

  For more information, see "Using Static IP Addresses" on page 27.

- For a file server that must always be assigned the same address, use Mac OS X Server's static IP mapping to always assign the same IP address to its Ethernet address.

  For more information, see "Assigning Static IP Addresses Using DHCP" on page 35.

- For DHCP configuration, set the LDAP options for DHCP clients. This automatically gives computers their needed directory information.

  For more information, see "Setting LDAP Options for a Subnet" on page 34.

- For client configuration on Mac OS X client computers, make sure the IPv4 configuration method in the Network pane of System Preferences is set to DHCP.

This configuration allows computers to be managed by an LDAP or Open Directory server, getting their network configuration information from DHCP. They can have access to truly static IP address or consistently assigned IP addresses on the same network. You also get centralized configuration for all computers.

Student Lab Configuration

The student lab configuration example is very much like the workgroup configuration example, but it adds NetBoot as an extra service that uses DHCP.

Along with DHCP providing centralized networking configuration, NetBoot standardizes startup environments by having each computer start up from a disk image on a central NetBoot server.

The configuration would be like the workgroup configuration example, with the following differences:

- There might be static-address resources.

  This depends on the lab composition. 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 take a server and printer to each class.

- NetBoot must be enabled and configured, along with 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 further configuration is necessary, and computers can be repurposed easily, because the hard disk can remain unchangeable.

With this configuration, computers on the network can be managed with an LDAP or Open Directory server, getting their network configuration information from DHCP. The computing environment is also centrally configured for all computers. New computers can be added or swapped out with minimal effort.

Coffee Shop Configuration

The coffee shop configuration is an example configuration for a dynamic addressing environment, one that requires no user management and provided no services other than web access, DNS access, or other service.

This example is characterized by lots of mobile users who pass through, use the Internet access, and move on.

This configuration can easily be used in situations like a college-commons wireless network or a wired courtesy office for visiting consultants.

> ***WARNING:*** If you host temporary unauthenticated users, make sure sensitive information on your LAN is protected behind a firewall on another network.

To use DHCP in this setting, you must have a working firewall configured for web access outbound traffic and DNS outbound lookups only. You might need to place this network outside your firewall and make sure the DHCP allocated IP addresses' network traffic is strictly controlled and monitored.

For more information, see Chapter 4, "Working with Firewall Service."

In this example, you might want to configure the DHCP service like this:

- **Make networking configuration automatic.** Set DHCP clients to get network configuration through DHCP.
- **Don't set options that clients shouldn't have.** Don't give DHCP clients more information about your organization than necessary using LDAP. You might want to configure Windows clients to have more network options.

  For more information, see "Setting WINS Options for a Subnet" on page 34.
- **Limit resource use.** Having many users on a subnet can lead to a lot of bandwidth use, so reduce the number of DHCP clients that can be connected simultaneously by restricting the number of addresses to be allocated.

  For more information, see "Creating Subnets in DHCP Service" on page 29.

- **Keep address turnover high.** Make the lease times on addresses as short as practical. This way, as users come and go, the addresses can be quickly reallocated.

  For more information, see "Creating Subnets in DHCP Service" on page 29.

- **Monitor your traffic.** Keep a close eye on DHCP connections and clients, firewall rule packet logging, or other monitoring tools. Open access points can be a liability if they are not guarded vigilantly.

Configuring DHCP to Use an Extra LDAP Server URL

The Server Admin application's DHCP module enables administrators to specify a single LDAP server URL for each subnet. If you want to specify multiple LDAP server URLs, you can edit the /etc/bootpd.plist file or use the serveradmin command-line tool (from a Terminal window).

Editing the /etc/bootpd.plist file to add multiple LDAP server URLs

After you create a subnet using Server Admin DHCP and specify a single LDAP server URL, you can inspect and modify the settings by editing the /etc/bootpd.plist file:

1. Open the /etc/bootpd.plist file in an editor.
2. Locate the tag `<string>` in between the tag `<array>` of the `dhcp_ldap_url` key.

```
<key>dhcp_ldap_url</key>
<array>
<string>ldap://server.example.com/dc=server,dc=example,dc=com</string>
</array>
```

3. Add another LDAP server URL by inserting a `<string>` tag below the existing `<string>` tag and entering your LDAP server URL between the open `<string>` and closed `</string>` tags.

```
<key>dhcp_ldap_url</key>
<array>
<string>ldap://server.example.com/dc=server,dc=example,dc=com</string>
<string>ldap://server2.example.com/dc=server2,dc=example,dc=com</string>
</array>
```

4. Save the bootpd.plist file and exit your editor.
5. If DHCP is running, restart the DHCP service so it can pick up the revised configuration.

   Using the Terminal you would enter:

```
$ sudo serveradmin stop DHCP
$ sudo serveradmin start DHCP
```

Using serveradmin to multiple LDAP server URLs

After you create a subnet using Server Admin DHCP and specify a single LDAP server URL, you can inspect and modify the settings using serveradmin. Do the following.

1. Inspect all DHCP subnet settings in the Terminal by entering:

```
$ sudo serveradmin settings dhcp:subnets
```

   Example result (excerpt):

...

```
dhcp:subnets:_array_id:498D8E6D-88A8-4048-8B3C-
     14D96F317447:dhcp_ldap_url:_array_index:0 = "http://ldapxxx:123/
     basename1"
```

...

2 Prepare a file with the serveradmin commands to add a second LDAP Server URL.

Because the individual elements of the dhcp_ldap_url array are not individually accessible, you cannot use the serveradmin create/delete idiom.

Example file contents:

```
dhcp:subnets:_array_id:498D8E6D-88A8-4048-8B3C-
     14D96F317447:dhcp_ldap_url:_array_index:0 = "http://ldapxxx:123/
     basename1"
dhcp:subnets:_array_id:498D8E6D-88A8-4048-8B3C-
     14D96F317447:dhcp_ldap_url:_array_index:1 = "http://ldapyyy:234/
     basename2"
```

Note: The array indexes start with 0. The old URL entry must be present even though you are just adding a second one. The entries must be in order.

3 Use the serveradmin tool to apply the settings from the file by entering:

```
$ sudo serveradmin settings < filename
```

Example result (the settings are confirmed):

```
dhcp:subnets:_array_id:498D8E6D-88A8-4048-8B3C-
     14D96F317447:dhcp_ldap_url:_array_index:0 = "http://ldapxxx:123/
     basename1"
dhcp:subnets:_array_id:498D8E6D-88A8-4048-8B3C-
     14D96F317447:dhcp_ldap_url:_array_index:1 = "http://ldapyyy:234/
     basename2"
```

4 If DHCP is running, restart the DHCP service so it can pick up the revised configuration by entering:

```
$ sudo serveradmin stop DHCP
$ sudo serveradmin start DHCP
```

DHCP Service for Mac OS X Clients Using DHCP with a Manual Address

The DHCP section of Server Admin permits each subnet address range to be enabled or disabled. When the subnet is enabled, the DHCP server allocates addresses in its range and dispenses other network information to clients that are configured as "Using DHCP."

When the subnet is disabled, the DHCP server does not allocate addresses from the subnet address range pool, but it does dispense other network information (such as DNS and LDAP server addresses) to clients that are configured as "Using DHCP with manual address" (static maps), as long as the client address is in the subnet range.

Enabling and disabling the subnet disables automatic address allocation for the address range, and *not* DHCP server responses to clients whose address is in the subnet range.

Where to Find More Information

Request for Comments (RFC) documents provide an overview of a protocol or service and explain how the protocol should behave.

If you're a novice server administrator, you'll probably find the background information in an RFC helpful.

If you're an experienced server administrator, you can find 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 about advanced configuration options, see the `bootpd` man page.

Working with DNS Service

3

This chapter describes how to set up, secure, and manage DNS service on your network.

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, which is 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 need to configure your network preferences with the IP address of the name server, which is provided by your ISP.

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

- You can't use DNS from your ISP or other source.
- You plan on making frequent changes to the name space 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 because your network's computer names and addresses are accessible to an outside organization (your ISP).

Mac OS X Server uses Berkeley Internet Name Domain (BIND) v9.4.1 for its implementation of DNS protocols. BIND is an open source implementation and is used by most name servers on the Internet.

About DNS Zones

Zones are the basic organizational unit of DNS. Zones contain records and are defined by how they acquire those records and how they respond to DNS requests.

There are three basic zones:

- Primary
- Secondary
- Forward

There are also several other kinds of zones not covered here.

Primary Zones

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

Secondary Zones

A secondary zone is a copy of a primary zone and is stored on a secondary name server. It has the following characteristics:

- Each secondary zone has a list of primary servers that it contacts for updates to records in the primary zone. Secondaries must be configured to request the copy of the primary zone 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 make sure that lookup requests are answered if the primary name server is down.

Secondary zones also have a refresh interval. This interval determines how often the secondary zone checks for changes from the primary zone. You can change the zone refresh interval by using BIND's configuration file. For more information, see the BIND documentation.

Forward Zones

A forward zone directs 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.

Forward zones also cache responses to queries they pass on. This can improve the performance of lookups by clients that use the forward zone.

Server Admin does not support creation or modification of a forward zone. To create a forward zone, you must configure BIND manually at the command line. For details, see the BIND documentation.

About DNS Machine Records

Each zone contains a number of records. These records are requested when a computer translates a domain name (like www.example.com) to an IP number. Web browsers, mail clients, and other network applications rely on zone records to contact the correct server.

Primary zone records are queried by others across the Internet so they can connect to your network services.

There are several kinds of DNS records available for configuration by Server Admin:

- **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 used for mail in a zone.
- **Name Server (NS):** Stores the authoritative name server for a zone.
- **Pointer (PTR):** Stores the domain name of an IP address (reverse lookup).
- **Text (TXT):** Stores a text string as a response to a DNS query.
- **Service (SRV):** Stores information about the services a computer provides.
- **Hardware Info (HINFO):** Stores information about a computer's hardware and software.

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

If you need access to other kinds of records, you must edit the BIND configuration files manually. For details, see the BIND documentation.

About Bonjour

With Bonjour, you can share nearly anything, including files, media, printers, and other devices, in innovative and easier ways. It simplifies traditional network-based activities like file sharing and printing by providing dynamic discoverability of file servers and Bonjour-enabled network printers.

Bonjour begins by simplifying the otherwise complex process of configuring devices for a network. To communicate with other devices using IP, a device needs special information like an IP address, a subnet mask, DNS addresses, a DNS name, and preconfigured search paths. Understanding these cryptic details and performing the subsequent configuration can be daunting for the average user.

When a new computer or device is added to a network by means of autoconfiguration, like a DHCP server, Bonjour configures the device using a technique called link-local addressing. (If a DHCP server is available, Bonjour uses the assigned IP address.)

With link-local addressing, the computer randomly selects an IP address from a predefined range of addresses set aside by the Internet Assigned Numbers Authority (IANA) for link-local addressing and assigns that address to itself. Addresses are in the range 169.254.xxx.xxx.

The device then sends a message over the network to determine whether another device is using the address. If the address is in use, the device randomly selects addresses until it finds one that is available. When the device has assigned itself an IP address, it can send and receive IP traffic on the network.

Before You Set Up DNS Service

This section contains information to consider before setting up DNS on your network. Because the issues involved with DNS administration are complex and numerous, do not set up DNS service on your network unless you're an experienced DNS administrator.

A good source of information about DNS is *DNS and BIND, 5th edition*, by Paul Albitz and Cricket Liu (O'Reilly and Associates, 2006).

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

Consider creating a mail alias, such as "hostmaster," that receives mail and delivers it to the person that runs the DNS server at your site. This permits users and other DNS administrators to contact you regarding DNS problems.

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

After a name server is provided with the name/address pair of a host in another domain (outside the domain it serves), the information is cached, ensuring 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 value for a domain name/IP address pair has expired, the entry is deleted from the name server's cache and your server requests the information 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 Gateway 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 IANA. IANA registration makes sure that domain names are unique across the Internet. (For more information, see www.iana.org.)

If you don't register your domain name, your network can't communicate over the Internet.

After you register a domain name, you can create subdomains as long as you set up a DNS server on your network to track 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 tracks information for its subdomains, such as host (computer) names, static IP addresses, aliases, and mail exchangers.

If your ISP handles your DNS service, you must inform them of changes you make to your domain name, including added subdomains.

The range of IP addresses used with a domain must be clearly defined before setup. These addresses are used exclusively for one specific domain, never by another domain or subdomain. Coordinate the range of addresses with your network administrator or ISP.

Step 2: Learn and plan

If you're new to DNS, learn and understand DNS concepts, tools, and features of Mac OS X Server and BIND. See "Where to Find More Information" on page 75.

When you're ready, plan your DNS Service. Consider the following questions:

- Do you need a local DNS server? Does your ISP provide DNS service? Can you use multicast DNS names instead?
- How many servers do you need? How many additional servers do you need for backup DNS purposes? For example, should you designate a second or third computer for DNS service backup?
- What is your security strategy to deal with unauthorized use?
- How often should you schedule periodic inspections or tests of DNS records to verify data integrity?
- How many services or devices (such as intranet websites or network printers) need a name?

There are two ways to configure DNS service on Mac OS X Server:

- **Use Server Admin.** This is the recommended method. For instructions, see "Setting Up DNS Service" on page 52.
- **Edit 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 must change the configuration file and the zone file. The configuration file is /etc/named.conf.

  The zone file name is based on the name of the zone. For example, the zone file example.com is /var/named/example.com.zone.

  If you edit named.conf to configure BIND, don't change the `inet` settings of the controls statement. Otherwise, Server Admin can't 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"; };
};
```

Important: In Mac OS X Server v0.5, the configuration and zone files used by Server Admin have changed. If you edit the named.conf and any zone files manually from Terminal, the information is used by DNS. However, the information does not appear in the DNS zones pane of Server Admin. Also, changes made in Server Admin are not made to the named.conf file.

Step 3: Turn DNS service on

Before configuring DNS service, turn on DNS. See "Turning DNS Service On" on page 51.

Step 4: Create a DNS zone and add machine records
Use Server Admin to set up DNS zones. See "Configuring Zone Settings" on page 53. After adding a primary zone, Server Admin creates a name server record with the same name as the Source of Authority (SOA).

For each zone you create, Mac OS X Server creates a reverse lookup zone. Reverse lookup zones translate IP addresses to domain names. (Compare with normal lookups, which translate domain names to IP addresses.)

Use Server Admin to add records to your zone. Create an Address record for every computer or device (such as a printer or file server) that has a static IP address and needs a name. Various DNS zone records are created from DNS machine entries.

Step 5: Configure secondary zones
If necessary, use Server Admin to configure secondary zones. See "Configuring Secondary Zone Settings" on page 54.

Step 6: Configure Bonjour
Use Server Admin to configure Bonjour settings. See "Configuring Bonjour Settings" on page 55.

Step 7: Configure logging
Use Server Admin to specify the information that gets logged by DNS service and the location of the log file. See "Changing DNS Log Detail Levels" on page 58.

Step 8: (Optional) Set up a mail exchange (MX) record
If you provide mail service over the Internet, set up an MX record for your server. See "Configuring DNS for Mail Service" on page 70.

Step 9: Configure your firewall
Configure your firewall to make sure your DNS service is protected from attack and accessible to your clients. See Chapter 4, "Working with Firewall Service."

Step 10: Start DNS service
Mac OS X Server includes a simple interface for starting and stopping DNS service. See "Starting DNS Service" on page 56.

Turning DNS Service On

Before you can configure DNS settings, turn on DNS service in Server Admin.

To turn DNS service on:

1. Open Server Admin and connect to the server.
2. Click Settings.
3. Click Services.

4 Select the DNS checkbox.

5 Click Save.

Upgrading DNS Configuration

Mac OS X Server v10.5 has been modified to manage DNS entries more efficiently. To take advantage of this, DNS records created on prior versions of Mac OS X Server must be upgraded.

After you upgrade to Mac OS X Server v10.5 and turn on DNS in Server Admin, the upgrade pane appears the first time you click DNS. (The upgrade pane appears only if you upgraded to Mac OS X Server v10.5. It will not appear if Mac OS X Server v10.5 was newly installed.)

The upgrade pane has two options:

- **Don't Upgrade:** If you choose to not upgrade your configuration, you cannot use Server Admin to automatically configure DNS. You can manually configure files using the /etc/named.conf file for DNS configuration and the /var/named file for Zone configuration.
- **Upgrade:** The Upgrade option converts DNS file records and then allows access to the DNS panes of Server Admin.

  When upgrading, backup files are created. If the files must be restored, they can be restored manually. Backup files are saved in the same folders where the original files are located.

Setting Up DNS Service

Set up DNS service by configuring the following three groups of settings in Server Admin:

- **Zones.** Use to configure a primary zone and computers that are part of the zone and to configure a copy of a primary zone stored on a secondary name server. This also sets information that determines if you permit zone transfers.
- **Bonjour.** Use to configure automatic Bonjour browsing.
- **Settings.** Use to configure and manage logs for DNS service and to set recursion for DNS service.

The following sections describe how to configure these settings. A final section explains how to start DNS service when you finish.

Configuring Zone Settings

Use Server Admin on the cluster controller to create a local DNS zone file and add records to it to map cluster nodes to corresponding IP addresses. The Open Directory service on the cluster controller requires this information to enable automatic server setup.

Important: In Mac OS X Server v10.5, the configuration and zone files used by Server Admin have changed. If you edit the named.conf and zone files manually from Terminal, the information is used by DNS. However, the information does not appear in the DNS zones pane of Server Admin. Also, changes made in Server Admin are not made to the named.conf file. It is recommended that you use Server Admin.

To configure DNS service zone settings:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Click Add Zone, then choose "Add Primary Zone (Master)."
6. Select the new zone.
7. In the Primary Zone Name field, enter the zone name.

   This is the fully qualified domain name of the primary server.
8. Enter the mail address of the zone's administrator.
9. Select "Allows zone transfer" to permit secondary zones to get copies of the primary zone data.
10. Add name servers for this zone by clicking the Add (+) button and entering the name in the Name Servers field.
11. Add mail exchangers for this zone by clicking the Add (+) button and entering the name in the Mail Exchangers field.

    This field is the basis for the computer's MX record.
12. Specify a mail server precedence number in the Priority field.

    Delivering mail servers try to deliver mail at lower numbered mail servers first. For more information, see "Configuring DNS for Mail Service" on page 70.
13. Click Expiration and enter the number of hours for each setting.

    Enter the amount of time the zone is valid. This is the zone's time to live (TTL) value. It determines how long any query response information can remain cached in remote DNS systems before requerying the authoritative server.

Enter the interval of time that the secondary zones should refresh from the primary zone.

Enter the interval of time between each retry if the refresh of the secondary zone fails.

Enter the amount of time after refreshing before the zone data expires.

14 Click Add Record, then choose "Add Alias (CNAME)."

To see a list of records for a zone, click the triangle to the left of the zone.

15 Select newAlias listed under the primary zone, then enter the alias information.

In the Alias Name field, enter the alias name. This field is the basis for CNAME records of the computer. Reverse lookup Pointer records are automatically created for the computer.

The fully qualified domain name of the computer appears beneath the alias name.

Add as many aliases as you want.

16 Click Add Record, then choose "Add Machines (A)."

17 Select newMachine listed under the primary zone, then enter the following machine information.

In the Machine Name field enter the hostname of the computer. This field is the basis for the A record of the computer. Reverse lookup Pointer records are automatically created for the computer.

Click the Add (+) button, then enter the IP address of the computer.

Enter any information about the hardware and software of the computer in the relevant text boxes. These are the basis for the HINFO record of the computer.

Enter any comments about the computer in the Comment text box. This field is the basis for the TXT record of the computer. You can store almost any text string in the comments text box up to 255 ASCII characters. For example, you can 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 about the computer.

18 Click Add Record, then choose "Add Service (SRV)."

19 Under the primary zone select the Homepage listed, then enter the service information.

20 Click Save.

Configuring Secondary Zone Settings

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 for updates to records in the primary zone.

Secondary zones must be configured to request the copy of the primary zone 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. Make sure the primary server is correctly configured and that sure zone transfers are enabled on the primary serve; then open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Click Add Zone, then choose "Add Secondary Zone (Slave)."
6. Select the new zone.
7. In the Secondary Zone Name field, enter a zone name.

   The zone name is the same as the primary zone defined on the primary name server.
8. Below the Primary Zone Addresses list, click the Add (+) button.
9. Enter the IP addresses for each primary server in this secondary zone.
10. Click Save.

Configuring Bonjour Settings

With Bonjour, you can easily connect a computer or other electronic device to an existing wired or wireless Ethernet network or you can create instant networks of multiple devices without additional network configuration.

To configure DNS service Bonjour settings:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Bonjour.
5. To activate wide-area Bonjour browsing, select "Enable automatic client Bonjour browsing for domain," then enter the network address.

   All primary zones provide clients with this domain for Bonjour browsing.
6. Click Save.

Configuring DNS Settings

You use the Settings pane in DNS to set the detail level of the DNS service log. You might want a highly detailed log for debugging or a less detailed log that only shows critical warnings.

You set recursive queries, which the DNS server fully answers the query (or gives an error). If the query is unanswered, it is forwarded to the IP addresses that you add in the Forwarder IP Addresses list.

To configure DNS settings:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Settings.
5. From the Log Level pop-up menu, choose the detail level as follows:
   - Choose Critical to record only critical errors, such as hardware errors.
   - Choose Error to record all errors not including warning messages.
   - Choose Warning to record all warnings and errors.
   - Choose Notice to record only important messages, warnings, and errors.
   - Choose Information to record most messages.
   - Choose Debug to record all messages.

   The log location is /Library/Logs/.
6. Below the "Accept recursive queries from the following networks" list, click the Add (+) button to add networks that recursive queries are accepted from, then enter the network address in the list.
7. Below the "Forwarder IP Addresses" list, click the Add (+) button to add networks that unauthorized queries get forwarded to, then enter the network address in the list.
8. Click Save.

Starting DNS Service

Use Server Admin to start DNS service.

Remember to restart the DNS service when you make changes to the DNS service in Server Admin.

To start DNS service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Start DNS (below the Servers list).

The service can take a few seconds to start.

Managing DNS Service

This section describes typical tasks you might perform after you set up DNS service on your server. Initial setup information appears in "Setting Up DNS Service" on page 52.

More advanced features require configuring BIND from the command-line and are not covered here.

You might want to monitor DNS status to:

- Troubleshoot name resolution problems.
- Verify how often the DNS Service is used.
- Look for unauthorized or malicious DNS service use.

This section discusses the following common monitoring tasks for DNS service.

Checking DNS Service Status

You can use Server Admin to check the status of DNS service.

To check DNS service status:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Overview to see whether the service is running, when it was started, and the number of zones allocated.

5 Click Log to review the service log.

Use the Filter field above the log to search for specific entries.

Viewing DNS Service Logs

DNS service creates entries in the system log for error and alert messages. The log view file is named.log. You can filter the log to narrow the number of viewable log entries and make it easier to find those you want to see.

To view logs:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Log and use the Filter field above the log to search for specific entries.

Changing DNS Log Detail Levels

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

To change the log detail level:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Settings.
5. Choose the detail level from the Log Level pop-up menu as follows:
   - Choose Critical to record only critical errors, such as hardware errors.
   - Choose Error to record all errors not including warning messages.
   - Choose Warning to record all warnings and errors.
   - Choose Notice to record only important messages, warnings, and errors.
   - Choose Information to record most messages.
   - Choose Debug to record all messages.
6. Click Save.

Stopping DNS Service

Use Server Admin to stop DNS service.

To stop DNS service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.

4 Click Stop DNS (below the Servers list).

5 Click Stop Now.

The service may take a few seconds to stop.

Enabling or Disabling Zone Transfers

In DNS, zone data is replicated among authoritative DNS servers by means of zone transfers. Secondary DNS servers (secondaries) use zone transfers to acquire their data from primary DNS servers (primaries). You must enable zone transfers if you want to use secondaries.

To enable or disable zone transfers:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Zones.

5 Select the Primary zone you want to change.

6 Click General.

7 Select or deselect "Allows zone transfer" to permit secondary zones to get copies of the primary zone data.

8 Click Save.

Enabling Recursion

Recursion fully resolves domain names into IP addresses. Applications depend on the DNS server to perform this function. Other DNS servers that query your DNS servers don't need to perform the recursion.

To prevent malicious users from changing the primary zone's records (referred to as cache poisoning) and to prevent unauthorized use of the server for DNS service, you can restrict recursion. However, if you restrict your private network from recursion, your users can't use your DNS service to look up names outside of your zones.

Disable recursion only if:

- No clients are using this DNS server for name resolution.
- No servers are using it for forwarding.

To enable recursion:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Settings.

5 Click the Add (+) button below the "Accept recursive queries from the following networks" list.

6 Enter the IP addresses for the servers that DNS will accept recursive queries from.

You can also enter IP address ranges.

7 Click Save.

If you enable recursion, consider disabling it for external IP addresses but enabling it for LAN IP addresses by editing BIND's named.conf file. However, any edits you make to the named.conf file will not show up in the DNS section of Server Admin. You can completely disable recursion by removing all entries from the network list. For more information about BIND, see www.isc.org/sw/bind.

Managing DNS Zones

DNS zones are managed using Server Admin. The following topics describe how to manage and modify DNS zones.

- "Adding a Primary Zone" on page 60
- "Adding a Secondary Zone" on page 61
- "Adding a Forward Zone" on page 62
- "Changing a Zone" on page 62
- "Deleting a Zone" on page 62

Adding a Primary Zone

Use Server Admin to add a primary zone to your DNS server.

To add a primary zone:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Zones.

5 Click Add Zone, then choose "Add Primary Zone (Master)."

6 Select the new zone.

7 In the Primary Zone Name field, enter the zone name.

This is the fully qualified domain name of the primary server.

8 Enter the mail address of the zone's administrator.

9 Select "Allows zone transfer" to permit secondary zones to get copies of the primary zone data.

10 Add nameservers for this zone by clicking the Add (+) button and entering the name in the Nameservers field.

11 Add mail exchangers for this zone by clicking the Add (+) button and entering the name in the Mail Exchangers field.

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

12 In the Priority field specify a mail server precedence number.

Delivering mail servers try to deliver mail at lower numbered mail servers first. For more information, see "Configuring DNS for Mail Service" on page 70.

13 Click Expiration and enter the number of hours for each setting.

Enter the amount of time the zone is valid. This is the zone's time to live (TTL) setting. It determines how long query response information can remain cached in remote DNS systems before requerying the authoritative server.

Enter the interval of time that the secondary zones should refresh from the primary zone.

Enter the interval of time between each retry if the refresh of the secondary zone fails.

Enter the amount of time after refreshing before the zone data expires.

14 Click Save.

Adding a Secondary Zone

Use Server Admin to add a secondary zone to your DNS server.

You perform the following steps on the secondary server. Before performing these steps, make sure the primary server is correctly configured and that zone transfers are enabled on the primary server.

To add a secondary zone:

1 On the secondary server, open Server Admin and connect to the primary server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Zones.

5 Click Add Zone, then click "Add Secondary Zone (Slave)."

6 Select the new zone.

7 In the Secondary Zone Name field, enter a zone name.

The zone name is the same as the primary zone defined on the primary name server.

8 Below the Primary Zone addresses list, click the Add (+) button.

9 Enter the IP addresses for each primary server in the secondary zone.

10 Click Save.

Adding a Forward Zone

A forward zone directs all lookup requests to other DNS servers. The forward zone also caches previous lookup requests for enhanced speed.

Use Server Admin to add a forward zone to your DNS server.

To add a forward zone:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Settings.

5 Click the Add (+) button below the Forwarder IP Addresses list.

6 Enter the IP addresses for the master servers for the forward zone.

A forward zone directs all lookup requests to other DNS servers. The forward zone also caches previous lookup requests for enhanced speed.

7 Click Save.

Changing a Zone

Use Server Admin to change zone settings. You might need to change the administrator mail address or domain name of a zone.

To change a zone:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Zones.

5 Select the zone you want to change.

6 Change the zone information as needed.

7 Click Save.

Deleting a Zone

When you delete a zone, all records associated with it are deleted.

To delete a zone:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Select the zone you want to delete.
6. Click Remove below the Zones list.
7. Click Save.

Importing a BIND Zone File

You might already have a BIND zone file from a DNS server of another platform. If so, instead of entering the information in Server Admin manually, you can use the BIND 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 and the Terminal application

  Otherwise, use the Server Admin DNS tools

Important: In Mac OS X Server v10.5, the configuration and zone files used by Server Admin have changed. If you edit the named.conf and zone files manually from Terminal, the information is used by DNS. However, the information does not appear in the DNS zones pane of Server Admin. Also, changes made in Server Admin are not made to the named.conf file. It is recommended that you use Server Admin.

To import a zone file:

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

   You need root privileges to edit named.conf.

   For example, for zone xyz.com described in zone file db.xyz.com in working zone folder /var/named/, the zone directive might look 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. Confirm that the zone file is added to the /var/named/ working zone folder.
3. Restart the DNS service using Server Admin.

Managing DNS Records

Each zone contains a number of records that are requested when a client computer translates a domain name (like www.example.com) to an IP number.

Web browsers, mail clients, and other network applications rely on a zone's records to contact the correct server.

The following topics describe how to add, modify, and delete DNS records.

- "Adding an Alias Record to a DNS Zone" on page 64
- "Adding a Machine Record to a DNS Zone" on page 65
- "Adding a Service Record to a DNS Zone" on page 65
- "Changing a Record in a DNS Zone" on page 66
- "Deleting a Record from a DNS Zone" on page 66

Adding an Alias Record to a DNS Zone

You must add records for each computer the DNS primary zone has responsibility for. Do not add records for computers this zone doesn't control.

An alias record or canonical name (CNAME) record is used to create aliases that point to other names. If you want this computer to have more than one name, add alias records to the zone.

To add a DNS alias record:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Select the zone this record is to be added to.
6. Click Add Record, then choose Add Alias (CNAME).

   This adds the alias record to the zone.
7. Select newAlias listed under the zone, then enter the alias information.

   In the Alias Name field, enter the alias name. This field is the basis for CNAME records of the computer. Reverse lookup Pointer records are automatically created for the computer.

   To use the fully qualified domain name of the computer, select the Fully Qualified checkbox.
8. Click Save.

Add as many aliases as you want by adding additional alias records.

Adding a Machine Record to a DNS Zone

You must add records for each computer the DNS primary zone has responsibility for. Do not add records for computers this zone doesn't control.

A machine record or address (A) record is used to associate a domain name with an IP address. Therefore, there can be only one machine for each IP address because there can't be duplicate IP addresses in a zone.

To add a DNS machine record:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Select the zone this record is to be added to.
6. Click Add Record, then choose Add Machine (A).

   This adds the machine record to the zone.
7. Select newMachine listed under the zone, then enter the following machine information.

   In the Machine Name field enter the hostname of the computer. This field is the basis for the A record of the computer. Reverse lookup Pointer records are automatically created for the computer.

   Click the Add (+) button, then enter the IP address of the computer.

   Enter any information about the hardware and software of the computer in the relevant text boxes. These are the basis for the HINFO record of the computer.

   Enter any comments about the computer in the Comment text box.

   This field is the basis for the TXT record of the computer.

   You can store up to 255 ASCII characters in the comments text box. You can include the physical location of the computer (for example, Upstairs server closet B), the computer's owner (for example, John's Computer), or any other information about the computer.
8. Click Save.

Adding a Service Record to a DNS Zone

You must add records for each computer the DNS primary zone has responsibility for. Do not add records for computers this zone doesn't control.

Service (SRV) records are used to define the target host and port where DNS service will work.

To add a DNS service record:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Select the zone this record is to be added to.
6. Click Add Record, then choose Add Service (SRV).

   This adds the service record to the zone.
7. Under the zone select the Homepage listed, then enter the service information.
8. Click Save.

Changing a Record in a DNS Zone

If you change the namespace for the domain, you must update the DNS records as often as that namespace changes. Upgrading hardware or adding to a domain name might also require updating the DNS records.

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

To change a record:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.
5. Click the triangle to the left of the zone that has the computer record to be edited.

   The list of records appears.
6. Select the record to be edited and make changes in the fields below the list.
7. Click Save.

Deleting a Record from a DNS Zone

When a computer is no longer associated with a domain name or usable address, delete the associated records.

To delete a record:

1. Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select DNS.

4 Click Zones.

5 Click the triangle to the left of the zone that has the computer record to be deleted.

The list of records appears.

6 Select the record to be deleted and click Remove below the list.

7 Click Save.

Securing the DNS Server

DNS servers are targeted by malicious computer users (hackers). DNS servers are susceptible to several kinds of attacks. By taking extra precautions, you can prevent the problems and downtime associated with hackers.

Several kinds of security attacks are associated with DNS service:

- DNS spoofing
- Server mining
- DNS service profiling
- Denial of service (DoS)
- Service piggybacking

DNS Spoofing

DNS spoofing is adding false data to the DNS Server's cache. This enables hackers to:

- 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 users into giving their bank account numbers and passwords to the hacker.

  Also, a falsified mail record could enable a hacker to intercept mail sent to or from a domain. If the hacker then forward that mail to the correct mail server after copying the mail, this can go undetected.

- 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 and auditing DNS records regularly.

If exploits are found in the current version of BIND, the exploits are 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 addresses 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, specify which IP addresses have permission to request zone transfers (your secondary zone servers) and deny all others.

Zone transfers are accomplished over TCP on port 53. To limit zone transfers, block zone transfer requests from anyone but your secondary DNS servers.

To specify zone transfer IP addresses:

1. Create a firewall filter that permits only IP addresses that are inside your firewall to access TCP port 53.
2. Follow the instructions in "Configuring Advanced Firewall Rules" in Chapter 4, "Working with Firewall Service," using the following settings:
   - Packet: Allow
   - Port: 53
   - Protocol: TCP
   - Source IP: the IP address of your secondary DNS server
   - Destination IP: the IP address of your primary DNS server

DNS Service Profiling

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

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

To alter BIND's version response:

1. Open a command-line text editor (for example vi, emacs, or pico).
2. Open named.conf for editing.
3. To the options brackets of the configuration file, add the following:

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

4 Save named.conf.

Denial of Service (DoS)

This kind of attack is common and easy. A hacker sends so many service requests and queries that a server uses all 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 enables 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 "Configuring Advanced Firewall Rules" on page 93. 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 done not so much by malicious intruders but by common Internet users who learn the trick from other users. They might feel that the DNS response time with their own ISP is too slow, so they configure their computer to query another DNS server instead of their own ISP's DNS servers. Effectively, there are more users accessing the DNS server than were planned for.

You can guard against this type of attack 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 don't provide this service to Internet users.

To prevent recursion entirely, see "Enabling Recursion" on page 59.

The most common balance is permitting recursion for requests coming from IP addresses in your own range but denying recursion to external addresses.

BIND enables 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];
               };
     };
```

For more information, see the BIND documentation.

Wide Area Bonjour Service Administration

If you computer or devices supports Bonjour it will automatically broadcast and discover services from other computers or devices using Bonjour. You can quickly and and easily network computers and devices that support Bonjour.

Bonjour browsing can be managed and secured though Server Admin. You can manage Bonjour browsing by designating a Bonjour domain for all computers and devices that use Bonjour. When you enable this feature all primary zones provide the designated Bonjour browsing domain to client computers to browse for Bonjour services.

To enable Bonjour browsing:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Bonjour.
5. Select the "Enable automatic client Bonjour browsing for domain" checkbox and enter a domain name for Bonjour browsing (for example, *bonjour.company.com*).
6. Click Save.

Common Network Administration Tasks That Use DNS Service

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

Configuring DNS for Mail Service

To provide mail service on your network, you must set up DNS so that incoming mail is sent to the correct mail host on your network.

When you set up mail service, you define a series of hosts, known as *mail exchangers* or *MX hosts,* each of which has a defined priority level. 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.

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

user-name@reliable.example.com

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 you want handled by your mail server and direct the mail to the correct computer.

When you set up an MX record, include a list of potential 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 precedence number (its priority). 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 receives the mail, it holds the mail and sends it to the main mail server when the main server becomes available, and then the main mail server delivers the mail.

Following is an example of an MX record that includes three computers that can receive mail for the example.com domain:

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 process above 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 accepts the mail.

Configuring DNS for mail service involves creating MX records in DNS for your mail servers. If your ISP provides DNS service, contact the ISP so they can enable your MX records. Follow these steps only if you provide your own DNS Service.

You might want to set up multiple servers for redundancy. If so, create an MX record for each auxiliary server.

To enable MX records for your mail server:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DNS.
4. Click Zones.

5. Select the zone this record is to be added to.
6. Click the triangle to the left of the zone.

   The list of records appear.
7. Click Add Record, then choose Add Machine (A).

   This adds a machine record to the zone.
8. In the Machine Name field, enter the hostname of the computer.

   If you want to use the fully qualified name of the computer, select the Fully Qualified checkbox and enter the fully qualified domain name of the computer.

   This field is the basis for the A record of the computer. Reverse lookup Pointer records are automatically created for the computer.
9. Click the Add (+) button and enter the IP addresses for the computer.
10. In the relevant text boxes, enter information about the hardware and software of the computer.
11. In the Comment text box, enter comments about the computer.

    This field is the basis for the TXT record of the computer.

    You can store up to 255 ASCII characters in the comments text box. You can include the physical location of the computer (for example, Upstairs server closet B), the computer's owner (for example, John's Computer), or any other information about the computer.
12. Click Save.
13. To add other names that you want this computer to have, click Add Record and choose Add Alias (CNAME).

    Add as many aliases as you want for your server.
14. In the Alias Name field, enter the alternate name for your computer.

    If you want to use the fully qualified name for the Alias, select the Fully Qualified checkbox and enter the fully qualified domain name.

    This field is the basis for the CNAME records of the computer. Reverse lookup Pointer records are automatically created for the computer.
15. In the Destination field, enter the computer name you are creating the alias for.

    If you want to use the fully qualified name for the Destination, select the Fully Qualified checkbox and enter the fully qualified domain name.
16. Click Save.
17. From the expanded Servers list, select Mail.
18. Click Settings, then click Advanced.
19. Click Hosting.

20 Click the Add (+) button.

21 In the Local Host Alias field, enter the alias name you created earlier.

22 Click OK, then click Save.

23 Repeat Steps 7 through 22 for each mail server.

24 Click Save.

Setting Up Namespace Behind a NAT Gateway

If you're behind a NAT gateway, you have a special set of IP addresses that are usable only in the NAT environment. If you were to assigned a domain name to these addresses outside the NAT gateway, none of the domain names would resolve to the correct computer. For more information about NAT, see Chapter 5, "Working with NAT Service," on page 111.

However, you can 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 permit resolution of names outside the routed area.

Your client 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. For more information, see "Linking a LAN to the Internet Through One IP Address" on page 118.

If you set up namespace behind the gateway, names entered by users outside the NAT gateway won't resolve to the addresses behind it. 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, see "Configuring Port Forwarding" on page 113.

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

Network Load Distribution (Round Robin)

BIND permits simple load distribution using an address-shuffling method known as *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 can't monitor current server load or processing power. It only 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. 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 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.

To mitigate the effects of local caching, you might want the zone's time-to-live (TTL) number to be fairly short.

Setting Up a Private TCP/IP Network

If you have a LAN that has a connection to the Internet, set up your server and computers with IP addresses and other information that are unique to the Internet. You obtain IP addresses from your ISP.

If your LAN will not connect to the Internet and you want to use TCP/IP for transmitting information about your network, set up a private TCP/IP network. When you set up a private network, you choose IP addresses from the following blocks of IP addresses that IANA 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, register with an Internet registry and use the IP addresses provided by the registry when setting up your private network. Otherwise, when you connect to the Internet, every computer on your network must be reconfigured.

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

Hosting Several Internet Services with a Single IP Address

You can have one server that supplies all Internet services (such as mail or web). These services can run on one computer with a single IP address.

You can have multiple host names in the same zone for a single server. For example, you might want to have the domain name www.example.com resolve to the same IP address as ftp.example.com or mail.example.com. This domain appears to be several servers to anyone accessing the services, but they are all one server at one IP address.

Setting up DNS records for this service is easy: add aliases to the machine DNS record. Setting up DNS names for these services does not enable or configure the services. It provides names that are easy to remember for each service offered. This can simplify 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.

  Be sure to select the mail server checkbox on the machine pane.
- 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 on FTP clients.

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

Hosting Multiple Domains on the Same Server

You can have one server that supplies all Internet services (such as mail or web) for several different domain names. For example, you can have the domain name www.example.com resolve to the same IP address as www.server.org. This domain appears to be several servers to anyone accessing the domains, but they are all one server at one IP address.

Setting up DNS records for this service is easy: add a DNS zone and then add your host names and server information to that zone.

Setting up the DNS names for these services does not enable or configure the service for these domain names. This configuration is used with virtual domain hosting in mail and web services.

Where to Find More Information

For More Information About DNS and BIND

See the following:

- *DNS and BIND, 5th edition*, by Paul Albitz and Cricket Liu (O'Reilly and Associates, 2006)
- The International Software Consortium website:

  www.isc.org and www.isc.org/sw/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 explain 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 technical details about a protocol in its RFC document.

You can search for RFC documents by number at www.ietf.org/rfc.html.

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

Working with Firewall Service

4

This chapter describes how to set up and manage Firewall service in Mac OS X Server.

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

Turning on Firewall service is similar to erecting a wall to limit access to your network. Firewall service scans incoming IP packets and rejects or accepts these packets based on rules you use to configure Firewall service.

You can restrict access to any IP service running on the server, and you can customize rules for incoming clients or for a range of client IP addresses.

About Firewall Service

Firewall service is configured using Server Admin. You can also configure certain settings by manually editing configuration files.

The illustration below shows an example firewall process.

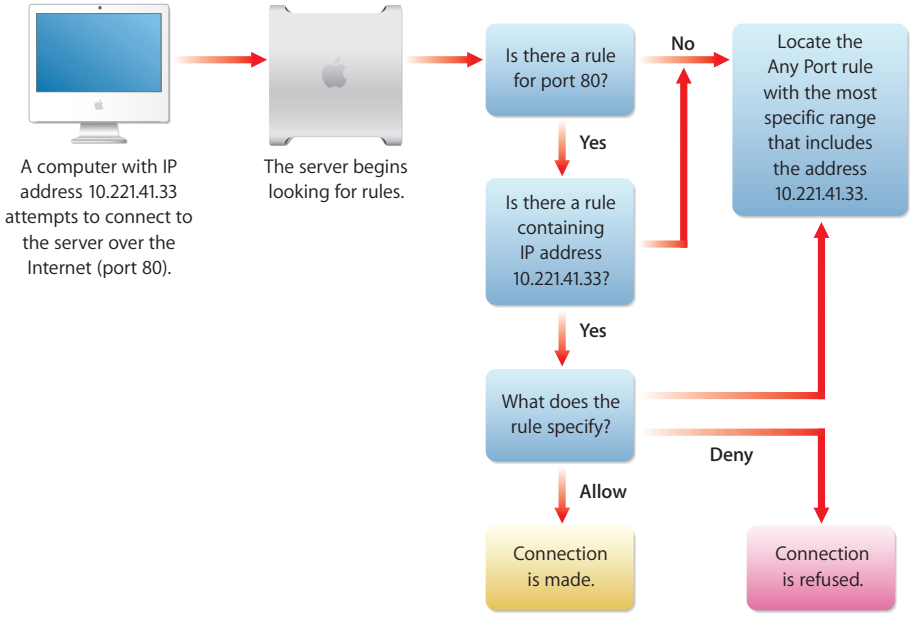

Services such as Web and FTP are identified on your server by a Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) port number. When a computer tries to connect to a service, Firewall service scans the rule list for a matching port number.

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

The rules you set are applied to TCP packets and 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. Other ports are dynamically opened to permit specific responses to queries initiated from the server. To permit remote access to other services on your computer, open more ports using the Services section of the Settings pane.

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 must 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 greater degree of control over their servers. For example, workgroups in a large business, or schools in a school system, can use Firewall service to control access to their own servers.

Firewall service 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. This permits packets that would otherwise be denied.

Basic Firewall Practices

By default, Mac OS X Server uses a simple model for a useful, secure firewall. If a firewall is too restrictive, the network behind it can be too isolated. If a firewall is too permissive, it fails to secure the assets behind it.

Adhering to the following aspects of the basic model provide maximum flexibility and utility with minimum unintended risk:

- Permit essential IP activity.

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

- Permit service-specific activity.

  Service-specific activity refers to network packets destined for specific service ports, such as web service or mail service. By permitting traffic to access ports with designated, configured services, you permit access through the firewall on a per-service basis.

  These services are expressed as medium-priority rules and correspond to check boxes in the Service pane of Firewall settings. You make these changes based on your settings and address groups.

- Deny packets not already permitted.

  This is the final catch-all practice. If a packet or traffic to a port is unsolicited, the packet or traffic is discarded and not permitted to reach its destination. This is expressed as low-priority (high-numbered) rules, visible in the Advanced pane of Firewall service 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 Server Admin, 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, Server Admin 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 as follows:

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

Otherwise, it disables the firewall as follows:

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

The rules loaded in the firewall remain regardless of this setting. They are ignored when the firewall is disabled.

Like most startup items, the IPFilter startup item opens in a predetermined order and only after prerequisite startup items have completed. In Mac OS X Server, the login window is presented while startup items can still be running. It is therefore feasible to log in before the firewall has activated its configured settings.

The startup item that sets up the firewall should generally finish a few minutes after starting the system.

About Firewall Rules

When you start Firewall service, the default configuration denies access to incoming packets from remote computers except through 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 permitted.

You can then add IP rules to permit 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 90.

What a Firewall Rule Is

A firewall rule is a set of characteristics for an IP packet, coupled with an action to be taken for each packet that matches the characteristics. The characteristics might include the protocol, source or destination address, source or destination port, or network interface.

Addresses might be expressed as a 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 the rule applies to, 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 computers in a company and the last segment might belong to a specific computer on one floor of a building.

Subnet Mask

A subnet mask indicates the segments in the specified IP address that can vary on a given network and by how much. The subnet mask is given in Classless InterDomain 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 that the first 16 bits (the first two sets of numbers separated by periods) are used to represent the network (so 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 the IP address followed by a colon (:) and the netmask. A netmask is a group of 4 numbers, each from 0 to 255, separated by periods equivalent to the slash in CIDR notation.

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	4.29×10^9
/2	192.0.0.0	2.14×10^9
/3	224.0.0.0	1.07×10^9
/4	240.0.0.0	5.36×10^8
/5	248.0.0.0	1.34×10^8
/6	252.0.0.0	6.71×10^7
/7	254.0.0.0	3.35×10^7
/8	255.0.0.0	1.67×10^7
/9	255.128.0.0	8.38×10^6
/10	255.192.0.0	4.19×10^6
/11	255.224.0.0	2.09×10^6
/12	255.240.0.0	1.04×10^6
/13	255.248.0.0	5.24×10^5
/14	255.252.0.0	2.62×10^5
/15	255.254.0.0	1.31×10^5
/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 permitted 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. You can change the range by changing the subnet mask.

When you indicate a range of potential 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	Example 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 Firewall Settings Services pane operate with the rules shown in the Advanced pane.

Usually, the broad rules in the Advanced pane block access for all ports. These are lower-priority (higher-numbered) rules and are applied after the rules in the Services pane.

The rules created with the Services pane open access to specific services and are higher priority. They take precedence over those created in the Advanced pane.

If you create multiple rules in the Advanced pane, the precedence for a rule is determined by the rule number. This number corresponds to the order of the rule in the Advanced pane.

Rules can be reordered by dragging them in the list in the Firewall Settings Advanced pane.

For most normal uses, opening access to designated services in the Advanced pane is sufficient. If necessary, you can add more rules using the Advanced pane.

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, the rules you create apply to all of those IP addresses.

Editing IPv6 Firewall Rules

When you configure and use the Firewall service in Server Admin, by default both `ipfw` and `ip6fw` are started. However, all IPv6 traffic except for local traffic is blocked.

You can override the IPv6 rules by using the `ip6fw` tool, but the after the Firewall service or the server is restarted your rules are overwritten.

You can control how Firewall in Server Admin manages the IPv6 firewall with the following two keys in the /etc/ipfilter/ip_address_groups.plist file:

```
<key>IPv6Mode</key>
<string>DenyAllExceptLocal</string>
<key>IPv6Control</key>
<true/>
```

The IPv6Mode key allows you to control which IPv6 rules are applied. There are three possible setting for the IPv6Mode string:

- DenyAllExceptLocal
- DenyAll
- NoRules

By default, the IPv6Mode key has the string set to DenyAllExceptLocal. This setting applies the following rules, which denies all IPv6 traffic but permits local network traffic:

```
add 1 allow udp from any to any 626
add 1000 allow all from any to any via lo0
add 1100 allow all from any to ff02::/16
65000 deny ipv6 from any to any
```

If you set the IPv6Mode string to DenyAll, only the following rule is applied, blocking all IPv6 traffic.

```
65000 deny ipv6 from any to any
```

If you set the IPv6Mode string to NoRules, no rules are created for IPv6. If your network is entirely IPv6, you may want to use this rule and use the `ip6fw` tool to create override rules for IPv6 and create a script that reapplies the rules when the Firewall services or server restarts.

The IPv6Control key allows you to set a boolean value that determines if `ip6fw` starts or stops when `ipfw` starts or stops. If the value is set to true, `ip6fw` starts and stops when `ipfw` start or stops. If the value is set to false, only `ipfw` starts or stops. By default the value is set to true.

Setup Overview

After you decide the types of rules to configure, follow the steps below to set up Firewall service. If you need more help to perform these steps, see "Setting Up Firewall Service" on page 87 and the other topics referred to in the steps.

Step 1: Learn and plan

If you're new to working with Firewall service, learn and understand firewall concepts, tools, and features of Mac OS X Server and BIND. For more information, see "About Firewall Rules" on page 80.

Then determine 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 are more likely to be restricted to your local subnet.

After you decide the services to protect using Firewall service, determine the IP addresses you want to permit access to your server and the IP addresses you want to deny access to your server. Then configure the suitable rules.

Step 2: Turn Firewall service on

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

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

Step 3: Configure Firewall Address Groups settings

Create an IP address group that the firewall rules will apply to. By default, an IP address group is created for all incoming IP addresses. Rules applied to this group affect all incoming network traffic. See "Configuring Address Groups Settings" on page 87.

Step 4: Configure Firewall Services settings

Activate service rules for each address group. In the Services pane, you can activate rules based on address groups as destination IP numbers. See "Configuring Services Settings" on page 88.

Step 5: Configure Firewall Logging settings
Use logging settings to enable Firewall service event logging. You can also set what types and how many packets get logged. See "Configuring Logging Settings" on page 89.

Step 6: Configure Firewall Advanced settings
Configure advanced firewall rules that are used to further configure all other services, strengthen your network security, and fine-tune your network traffic through the firewall. See "Configuring Advanced Firewall Rules" on page 93.

By default, all UDP traffic is blocked, except traffic arriving in response to an outgoing query. Apply rules to UDP ports sparingly, if at all, because denying some UDP responses could inhibit normal networking operations.

If you configure rules for UDP ports, don't select the "Log all allowed packets" option in the Firewall Logging settings pane in Server Admin. Because UDP is a connectionless protocol, every packet to a UDP port is logged if you select this option.

To learn how IP rules work, read "About Firewall Rules" on page 80.

Step 7: Turn Firewall service on
You turn Firewall service on using Server Admin. See "Starting Firewall Service" on page 89.

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

Turning Firewall Service On

Before you can configure Firewall settings, you must turn Firewall service on in Server Admin.

To turn Firewall service on:

1. Open Server Admin and connect to the server.
2. Click Settings.
3. Click Services.
4. Click the Firewall checkbox.
5. Click Save.

Setting Up Firewall Service

You set up the Firewall service by configuring the following settings on the Settings pane for Firewall service in Server Admin.

- **Address Groups:** Use to configure groups of IP addresses that firewall rules are applied to.
- **Services:** Use to configure which services are permitted to send and receive information through the firewall.
- **Logging:** Use to enable Firewall service event logging and set the type and number of packets that are recorded.
- **Advanced:** (Optional) Use to configure advanced rules and set rule precedence.

The following sections describe the tasks for configuring these settings. A final section tells you how to start the Firewall service after you configure it.

Configuring Address Groups Settings

You can define groups of IP addresses for your firewall rules. Then you can use these groups 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 addresses and subnet mask in CIDR notation (192.168.2.0/24), or IP addresses and subnet mask in netmask notation (192.168.2.0:255.255.255.0).

To configure address group settings:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Address Groups.
5. Below the Address Group pane, click the Add (+) button.
6. In the Group name field, enter a group name.
7. Enter the addresses and subnet mask you want the rules to affect.

   Use the Add (+) and Delete (–) buttons.

   To indicate any IP address, use the word "any."
8. Click OK.
9. Click Save.

Configuring Services Settings

By default, Firewall service permits all UDP connections and blocks incoming TCP connections on ports that are not essential for remote administration of the server. Also, by default, stateful rules are in place that permit specific responses to outgoing requests.

Before you turn on Firewall service, make sure you've set up rules permitting access from IP addresses you choose; otherwise, no one can access your server.

You can easily permit standard services through the firewall without advanced and extensive configuration. Standard services include:

- SSH access
- Web service
- Apple File service
- Windows File service
- FTP service
- Printer Sharing
- DNS/Multicast DNS
- ICMP Echo Reply (incoming pings)
- IGMP
- 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 affects connections already established with the server. For example, if you deny all access to your FTP server after starting firewall service, computers already connected to your FTP server are disconnected.

To configure the firewall standard services:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Services.
5. From the Edit Services for pop-up menu, select an address group.
6. For the address group, choose to permit all traffic from any port or to permit traffic on designated ports.
7. For each service you want the address group to use, select Allow.

If you don't see the service you need, add a port and description to the services list.

To create a custom rule, see "Configuring Advanced Settings" on page 89.

8 Click Save.

Configuring Logging Settings

You can choose the types of packets to log. You can log only the packets that are denied access, only the packets that are permitted access, or both.

Each logging option can generate many log entries but there are ways to limit the volume. You can:

- Log only permitted packets or denied packets, instead of all packets.
- Log packets only as long as necessary.
- Use the Logging Settings pane to limit the total number of packets.
- Add a count rule in the Advanced Settings pane to record the number of packets that match the characteristics you're interested in measuring.

To set up firewall logs:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Logging.

5 Select the "Enable logging" checkbox and choose to log permitted packets, denied packets, or a designated number of packets.

6 Click Save.

Configuring Advanced Settings

You use the Advanced Settings pane in Server Admin to configure specific rules for Firewall service. This is an optional configuration step for Firewall service.

For more information, see "Configuring Advanced Firewall Rules" on page 93.

Starting Firewall Service

By default, Firewall service blocks incoming TCP connections and denies UDP packets, except those received in response to outgoing requests from the server.

Before you turn on Firewall service, make sure you've set up rules permitting access from IP addresses you choose; otherwise, no one can access your server.

If you add or change a rule after starting Firewall service, the new rule affects connections already established with the server. For example, if you deny all access to your FTP server after starting Firewall service, computers already connected to your FTP server are disconnected.

To start Firewall service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click the Start Firewall button below the Servers list.

Managing Firewall Service

After you have set up the Firewall service, you can use Server Admin to perform day-to-day management tasks.

Stopping Firewall Service

You use Server Admin to stop Firewall service.

To stop Firewall service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Stop Firewall.

Creating an Address Group

Use Server Admin to create address groups for Firewall service.

To create an address group:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Address Groups.
5. Below the IP Address Groups list, click the Add (+) button.
6. In the Group name field, enter a group name.
7. Enter the addresses and subnet mask you want the rules to affect.

Use the Add (+) and Delete (–) buttons.

To indicate any IP address, use the word "any."

8 Click OK.

9 Click Save.

Editing or Deleting an Address Group

You can edit address groups to change the range of IP addresses affected. 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), IP address and network mask in CIDR notation (192.168.2.0/24), or IP address and network mask in netmask notation (192.168.2.0:255.255.255.0).

To edit or delete an address group:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Address Groups.

5 From the IP Address Groups list, select the group name.

6 Choose from the following:

To edit an IP address group, click the Edit (/) button below the list.

To delete an IP address group, click the Delete (–) button below the list.

7 Edit the Group name or addresses as needed and click OK.

8 Click Save.

Duplicating an Address Group

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

To duplicate an address group:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Address Groups.

5 From the IP Address Groups list, select the group name.

6 Below the IP Address Groups list, click the Duplicate button.

Adding to the Services List

You can add custom ports to the Services list. This enables you to open specific ports to your address groups without creating an advanced IP rule.

To add to the Services list:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Services.

5 Below the services list, click the Add (+) button.

6 Enter a rule name for the service.

7 Enter a single port (for example, 22) or a port range (for example, 650-750).

8 Choose a protocol.

If you want a protocol other than TCP or UDP, use the Advanced settings to create a custom rule.

9 Click OK

10 Click Save.

Editing or Deleting Items in the Services List

You can remove or edit the ports the Services list. This enables you to customize service choices for ease of configuration.

To change the Services list:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Services.

5 Select the service you want to change, then do the following:

To edit the service list, click the Edit (/) button below the services list.

To delete the service list, click the Delete (–) button below the services list.

6 Edit the name, port, or protocol as needed, and click OK.

7 Click Save.

Configuring Advanced Firewall Rules

You use the Advanced Settings pane in Server Admin to configure specific rules for Firewall service. Firewall rules contain originating and destination IP addresses with subnet masks. They also specify what to do with incoming network traffic. 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), IP address and subnet mask in CIDR notation (192.168.2.0/24), or IP address and subnet mask in netmask notation (192.168.2.0:255.255.255.0).

To set up an advanced firewall rule:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Advanced.
5. Click the Add (+) button.

   Alternatively, you can select a rule similar to the one you want to create, click Duplicate, and then click Edit.
6. In the Action pop-up menu, select whether this rule permits or denies access.

   If you choose Other, enter the action desired (for example, log).
7. From the Protocol pop-up menu, choose a protocol.

   If you choose Other, enter the protocol desired (for example, icmp, esp, ipencap).
8. From the Service pop-up menu, choose a service.

   To select a nonstandard service port, choose Other.
9. If desired, choose to log all packets that match the rule.
10. For the source of filtered traffic, choose an address group from the Address pop-up menu.

    If you don't want to use an existing address group, enter the source IP address range (using CIDR notation) you want to filter.

    If you want it to apply to any address, choose "any" from the pop-up menu.
11. If you selected a nonstandard service port, enter the source port number.
12. For the destination of filtered traffic, choose an address group from the Source pop-up menu.

    If you don't want to use an existing address group, enter the destination IP address range (using CIDR notation).

    If you want it to apply to any address, choose "any" from the pop-up menu.

13 If you selected a nonstandard service port, enter the destination port number.

14 From the Interface pop-up menu that this rule will apply to, choose In or Out.

"In" refers to the packets being sent to the server.

"Out" refers to the packets being sent from the server.

15 If you select Other, enter the interface name (en0, en1, fw1, and so on).

16 Click OK.

17 Click Save to apply the rule immediately.

Editing or Deleting Advanced Firewall Rules

You can remove or edit advanced firewall rules. If you think you'll use a rule again and only want to disable it, you can 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 computers are connected to your web server and you change the rule to deny all access to the server, connected computers are disconnected.

To change an advanced firewall rule:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Advanced.

5 Select the rule you want to change, then do the following:

To edit the services list, click the Edit (/) button below the services list.

To delete a rule, click the Delete (–) button below the services list.

6 Edit the rule as needed, and click OK.

7 Click Save.

Changing the Order of Advanced Firewall Rules

The priority level of an advanced firewall rule is determined by its order in the Advanced Rules list. Default rules that are locked cannot be reordered in the list.

To change the rule order:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Advanced.

5 Drag the rules to reorder them in the needed sequence.

6 Click Save.

Troubleshooting Advanced Firewall Rules

Advanced firewall configuration settings accept any input, assuming you are correctly configuring a rule.

Errors are not noticed until the rules are saved and Server Admin applies all rules using the `ipfw` command. Then, the first rule with a syntax error causes the operation to stop, and an error message is logged.

This error message does not indicate which rule is invalid, but all valid rules before the invalid one are loaded in the firewall.

The following section describes how you can determine which rule is invalid.

To determine which rule is invalid:

1 Read the error message in the log.

2 Wait a few minutes for Server Admin to show the active rules in the Firewall Overview pane.

3 Compare the list of active rules in the Firewall Overview pane with the rule list in the Settings section.

4 Inspect the contents of /etc/ipfilter/ipfw.conf.apple file to see which rules Server Admin tried to load in the firewall.

The first rule in the file that is not present in the Firewall Overview pane is likely the invalid one. However, there might be more invalid rules after that one.

5 If the rule corresponds to one from the Advanced Settings pane, disable it or correct it.

Disabled rules appear in the /etc/ipfilter/ipfw.conf.apple file preceded by a comment character so they are not processed by the `ipfw` tool.

Enabling Stealth Mode

You can hide your firewall by choosing not to send a connection failure notification to any connection that is blocked by the firewall. This is called stealth mode and it effectively hides your server's closed ports.

For example, if a network intruder tries to connect to your server, even if the port is blocked, he or she knows that there is a server and can find other ways to intrude.

If stealth mode is enabled, instead of being rejected, the hacker won't receive notification that an attempted connection took place.

To enable stealth mode:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Advanced.
5. Select "Enable for TCP," "Enable for UDP," or both, as needed.
6. Click Save.

Adaptive Firewall

Mac OS X v10.5 uses an adaptive firewall that dynamically generates a firewall rule if a user has 10 consecutive failed login attempts. The generated rule blocks the user's computer for 15 minutes, preventing the user from attempting to log in.

The adaptive firewall helps to prevent your computer from being attacked by unauthorized users. The adaptive firewall does not require configuration and is active when you turn on your firewall.

Resetting the Firewall to the Default Setting

A server can become unreachable for remote administration due to an error with the firewall configuration. In such a case, you must reset the firewall to its default state so Server Admin can access the server.

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

To reset the firewall to its default setting:

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

   ```
   $ ipfw -f flush
   ```
6. Edit the /etc/hostconfig file and set IPFILTER=-YES-.
7. Complete the startup sequence in the login window by entering `exit`:

   The computer starts up with the default firewall rules and firewall enabled. Then use Server Admin 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 (hackers). To maintain the security of your computers and user information, you must monitor firewall activity and deter potential threats. This section explains how to log and monitor your firewall.

Checking the Status of Firewall Service

Use Server Admin to check the status of Firewall service.

To check Firewall service status:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Overview to see whether the service is running, the number of active static and dynamic rules configured, the number of matching packets, and the number of bytes in matching packets handled by the firewall.

5 Click Log to review the Firewall service log.

6 To view a list of active firewall rules, click Active Rules.

This list includes the priority number for each rule, the number of matching packets and the number of bytes in matching packets, and a description of the rule.

Viewing Firewall Active Rules

Use Server Admin to view a simple summary of active firewall rules.
The Active Rules pane shows the number 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 rules are generated and saved in a file, and the `ipfw` command is invoked to load the rules into service.

As part of the flush operation, the number of packets and bytes associated with each rule are cleared.

The Active Rules pane provides a snapshot of the state of the firewall. When viewing this pane, dynamic rules might be shown with static rules.

Dynamic rules come and go in a matter of seconds, in response to network activity. They are the result of rules that include a keep-state clause (stateful rules). The Active Rules pane shows the rule number of the stateful rule that was triggered to create the dynamic rule.

To view active firewall rules:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Active Rules.

   A list of the rules appears, with a description of each rule in ipfw code format, the priority, packet count, and total bytes handled.

Viewing the Firewall Service Log

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

The log view shows the contents of /var/log/ipfw.log. You can refine the view using the text filter box.

To view the Firewall service log:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Log.

   To search for specific entries, use the Filter field above the log.

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 through 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 permit the remote client at 10.221.41.33:721 to access the server 192.168.12.12 on the LPR printing port 515 through 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 Network Address Translation (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. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Logging.
5. Make sure "Log all denied packets" is selected.
6. To view log entries, click Log.
7. In the text filter box, enter the word "unreach."

Viewing Packets Logged by Firewall Rules

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

To view filtered packets:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Logging.
5. Make sure "Log all allowed packets" is selected.

   If you have not turned on logging for a particular rule, see "Editing or Deleting Advanced Firewall Rules" on page 94.
6. To view log entries, click Log.

7 Enter the word "Accept" in the text filter box.

Practical Firewall Examples

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

Using Firewall with NAT

The firewall must be enabled to use NAT. Enabling NAT creates a divert rule in the firewall configuration.

Although Server Admin permits NAT service and Firewall service to be enabled and disabled independently, NAT service can operate only if both NAT and Firewall services are enabled. An essential part of NAT is the packet divert rule used in the firewall.

The firewall rule you set up instructs 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 must create or know the address group that corresponds to the LAN.

For detailed information about setting up a NAT LAN, see "Linking a LAN to the Internet Through One IP Address" on page 118.

Blocking Web Access to Internet Users

This section describes how you can permit users on your subnet to access your server's web service and deny access to the general public on the Internet.

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

To block web access using an advanced rule:

1 In Server Admin, create an address group named "LAN" with the address range 10.0.1.1/24.

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

For more information, see "Creating an Address Group" on page 90.

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

For more information, see "Configuring Advanced Firewall Rules" on page 93.

To block web access using standard rules:

1. In Server Admin, create an address group named "Web Server" with the address 10.0.2.1.

   For more information, see "Creating an Address Group" on page 90.
2. Click Settings, then click Services.
3. From the Edit Services for pop-up menu, choose the "Web Server" address group.
4. Select "Allow only traffic from 'Web Server' to these ports."
5. Select the HTTP- web service checkbox.
6. Click Save.

Logging Internet Access by Local Network Users

This section describes how you can permit users on your LAN to have access to another server's Web service and log their access attempts to the general public on the Internet.

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

To log Internet access by local network users:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Services.
5. From the Edit Services pop-up menu, choose the address group "any."
6. Select to allow traffic for the group "Web Server" on the designated web services port.
7. Select Allow Web Service.
8. Click Save.
9. Click Logging.
10. Select the "Enable logging" checkbox.
11. Select "Log all allowed packets."

    The logs are visible in the Log pane.

Blocking Junk Mail

This section describes how to reject mail from a junk mail sender with an IP address of 17.128.100.0 (for example) and accept all other Internet mail.

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

To prevent junk mail from being delivered to users:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Services.
5. From the "Edit Services for" pop-up menu, choose the "any" address group.
6. Select Allow only traffic from "any" to these ports.
7. Enable "Mail SMTP Standard" in the ports list.
8. Select Address Groups.
9. To create an address range, click the Add (+) button and enter a name for the address group in the Group name field.
10. To indicate the junk mail sender's address, enter 17.128.100.0 in the "Address in group" list by clicking the Add (+) button and entering the address.
11. Click OK.
12. Click Services
13. From the "Edit Services for" pop-up menu, choose the newly created address group.
14. Select Allow only traffic from "*newaddressgroupname*" to these ports.
15. To disable mail transfer, deselect "Mail SMTP Standard" checkbox in the ports list.
16. Click Save.

Permitting a Customer to Access the Apple File Server

This section describes how to permit a customer with an IP address of 10.221.41.33 (for example) to access an Apple file server.

To grant a customer access to the Apple file server:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings, then click Services.
5. From the Edit Services for pop-up menu, choose "any."
6. In the service pane, deselect "Apple File Service."
7. Select Address Groups.
8. To create an address range, click the Add (+) button.
9. Name the address group.

10 To indicate the customer's address, enter 10.221.41.33 in the address range field.

11 Click OK.

12 Click Services.

13 Select the newly created address group.

14 To enable file access, select "Apple File Service" in the service pane.

15 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 that can harm your data or abuse your network resources. This section describes 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 that is denied access, by default the server 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 who are trying to connect to the server. This is one type of DoS attack.

Important: DoS attacks are rare, so make these settings only if you think your server might be vulnerable to an attack. If you deny ICMP echo replies, services that use ping to locate network services can't detect your server.

To prevent ping DoS attacks:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select Firewall.

4 Click Settings, then click Services.

5 Select the "any" address group.

6 Deselect "ICMP Echo (ping) reply."

7 Click Save.

Controlling or Enabling Peer-to-Peer Network Usage

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

You can disable P2P networking by blocking incoming and outgoing traffic on the port number used by the P2P application. You must 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 limit P2P network usage to IP addresses behind the firewall. To do so, open the P2P port for your LAN interface but continue to block the port on the interface connected to the Internet (the WAN interface). To learn how to make a firewall rule, see "Configuring Advanced Firewall Rules" on page 93.

Controlling or Enabling Network Game Usage

Sometimes network administrators must control the use of network games. The games might use network bandwidth and resources improperly or disproportionately.

You can disable network gaming by blocking traffic incoming and outgoing on the port number used by the game. You must 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 limit network game usage to IP addresses behind the firewall. To do so, open the relevant port on your LAN interface but continue to block the port on the interface connected to the Internet (the WAN interface). Some games require a connection to a gaming service for play, so this might not be effective. To learn how to make a firewall rule, see "Configuring Advanced Firewall Rules" on page 93.

You can open the firewall to specific games, permitting network games to connect to other players and game services outside the firewall. To do this, open up the relevant port on your LAN and WAN interface. Some games require more than one port to be open. For networking details, consult the game's documentation. To learn how to make a firewall rule, see "Configuring Advanced Firewall Rules" on page 93.

Preventing Network Virus Propagation

A virus can quickly propagate through your network and infect your computers. For example, if a computer on your network becomes infected with a virus that computer can be used to propagate the virus through your entire network.

One common avenue that a virus uses to propagate through your network is by mail. You can prevent virus from propagating through mail by scanning mail with clamav and keeping your virus definitions updated.

You can prevent other avenues of propagation by only running services that you need, using good network topology, and using good passwords.

The most important method is to keep your network computers up-to-date. Your computer should be set to check for updates once or twice a week.

For more information about preventing network viruses see *Mac OS X Server Security Configuration*.

TCP and UDP 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 access rules. To view the RFCs referenced in the tables, see www.faqs.org/rfcs.

TCP port	Used for	Reference
7	echo	RFC 792
20	FTP data	RFC 959
21	FTP control	RFC 959
22	secure shell (SSH) Open Directory replica setup	
23	Telnet	RFC 854
25	SMTP (mail)	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 (with 3659)	
110	POP3 (mail)	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 (mail access)	RFC 2060
201-208	AppleTalk	

TCP port	Used for	Reference
311	Server Admin SSL, AppleShare IP remote web administration, Server Monitor, Server Admin (servermgrd), Workgroup Manager (DirectoryService)	
389	LDAP (directory)	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	WebObjects	
1099, 8043	Remote RMI and RMI/IIOP access to JBoss	
1220	QTSS Admin	
1694	IP Failover	
1723	PPTP VPN	RFC 2637
2049	NFS	
2399	FileMaker data access layer	
3004	iSync	

TCP port	Used for	Reference
3031	Program Linking, remote AppleEvents	
3283	Apple Remote Desktop 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	Apple Remote Desktop 2.0 database	
5900	Apple Remote Desktop 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 alternative (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

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

Where to Find More Information

For more information about accessing and implementing the features of `ipfw`, the tool that controls Firewall service, see the `ipfw` man page.

Request for Comments (RFC) documents provide an overview of a protocol or service and describe how the protocol should behave.

If you're a novice server administrator, you'll probably find the background information in an RFC helpful.

If you're an experienced server administrator, you can find all 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 Internet Assigned Number Authority (IANA) maintains a list of well known ports and TCP and UDP ports that have been assigned by the organization for various protocols. The list can be found at: www.iana.org/assignments/port-numbers.

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

Working with NAT Service

5

This chapter describes how to set up and manage NAT service in Mac OS X Server.

Network Address Translation (NAT) is a protocol you use to give multiple computers access to the Internet using only one assigned public or external IP address. NAT permits you to create a private network that accesses the Internet through a NAT router or gateway. NAT is sometimes referred to as IP masquerading.

The NAT router takes all traffic from your private network and remembers internal addresses that have made requests. When the NAT router receives a response to a request, it forwards it to the originating computer. Traffic that originates from the Internet does not reach 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, "Working with DHCP Service," on page 25.

Enabling NAT also creates a divert rule in the firewall configuration. Server Admin permits the NAT service and the Firewall service to be enabled and disabled independently. However for the NAT service to function, the NAT service and the Firewall service must 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.

NAT LAN Configuration Overview

To configure a network segment as a NAT LAN, you must complete several steps. Each is necessary to create a functioning private network behind a NAT gateway. A detailed example of the setup is found in "Linking a LAN to the Internet Through One IP Address" on page 118.

You can also configure NAT using the Gateway Setup Assistant, which configures each of these services and starts NAT. For more information, see "About the Gateway Setup Assistant" on page 17.

The following section provides an overview of the configuration process.

Step 1: Choose your NAT gateway and interface functions
You must locate the NAT gateway on 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 NAT LAN clients will get IP addresses
You can assign your own static IP address in 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
You assign your public IP address to the WAN port and you assign your internal gateway's address to the LAN port.

Step 4: Turn NAT service on
Before configuring NAT service, you must turn NAT on. See "Turning NAT Service On" on page 113.

Step 5: Configure NAT settings
Use the NAT settings to set the network interface. See "Configuring NAT Service" on page 113.

Step 6: Configure port forwarding settings
Use the Terminal application to direct incoming traffic to your NAT network to a specific IP address behind the NAT gateway. See "Configuring Port Forwarding" on page 113.

Step 7: Start NAT service
After you configure NAT, start the service to make it available. See "Starting and Stopping NAT Service" on page 116.

Step 8: Start Firewall service
For the NAT service to operate, you must enable both the NAT service and the Firewall service. See "Stopping Firewall Service" on page 90.

Step 9: (Conditional) Configure and start DHCP service
If clients will have their addresses dynamically assigned, configure DHCP and start it now. See Chapter 2, "Working with DHCP Service."

Turning NAT Service On

Before you can configure NAT settings, you must turn on NAT service in Server Admin.

To turn NAT service on:

1. Open Server Admin and connect to the server.
2. Click Settings.
3. Click Services.
4. Select the NAT checkbox.
5. Click Save.

Configuring NAT Service

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

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

To configure NAT service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select NAT.
4. Click Settings.
5. Select "IP Forwarding and Network Address Translation (NAT)."
6. From the "External network interface" pop-up menu, choose the network interface that connects to the Internet or external network.
7. Click Save.

Configuring Port Forwarding

You can direct traffic coming in to your NAT network to a specific IP address behind the NAT gateway. This is called *port forwarding*.

Port forwarding lets you set up computers on the internal network that handle incoming connections without exposing other computers to outside connections. For example, you could set up a web server behind the NAT service and forward 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 many ports to one computer.

Enabling port forwarding requires the use of the Terminal application and administrator access to root privileges through `sudo`.

You must also create a plist file. The contents of the plist file are 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 use a plist editor instead of a command-line text editor, alter the following procedure to suit.

To forward port traffic:

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

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

2. Using a Terminal editor, add the following block of XML text to /etc/nat/natd.plist before the two lines at the end the file (</dict> and </plist>), substituting your 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.
4. Enter the following commands in the terminal:

```
$ sudo systemstarter stop nat
$ sudo systemstarter start nat
```

5. Verify that your changes remain by inspecting the /etc/nat/natd.conf.apple file.

The changes made, except for comments and those settings that Server Admin can change, are used by server configuration tools (Server Admin, Gateway Setup Assistant, and serveradmin).

6 Configure NAT service in Server Admin as needed.

For more information, see "Configuring NAT Service" on page 113.

7 Click Save.

8 Start NAT service.

Port Forwarding Examples

You can forward one or more ports to an IP address. The ports on the WAN side do not need to be the same as the ports on the LAN side, but they 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 17.128.128.128 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/nat/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 17.128.128.128 to corresponding ports on the private LAN address of 192.168.1.1.

The block of text to add to the /etc/nat/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>
```

Testing Port Forwarding Rules

After you configure your port forwarding rules you can test them by accessing the service from the public IP address of your NAT router. If you successfully access the services you have properly configured and tested your port forwarding rule.

For example, if you have a website hosted on a computer with the private IP address of 192.168.1.10 and your NAT router has a public IP address of 219.156.13.13 and a port forwarding rule that forwards port 80 to IP address 192.168.1.10, you would access the website by entering the public IP address (http://219.156.13.13) into your web browser.

If your port forwarding rules are correct, you will be forwarded to the computer that is hosting the website (192.168.1.10).

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 start DHCP on the NAT interface, so you must manage LAN addressing separately.

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

For the NAT service to operate, you must enable the NAT service and the Firewall service. For more information, see "Stopping Firewall Service" on page 90.

To start NAT service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select NAT.
4. Click the Start NAT button below the Servers list.

   When the service is running, the Stop NAT button is available.

Creating a Gateway Without NAT

You can use a computer as a gateway between network segments without translating IP addresses between public and private ranges. This is called *IP address forwarding*. Mac OS X Server supports IP address forwarding and can be configured using Server Admin.

You can have various network configurations that would use a gateway without NAT. For example, a server might be translating private IP addresses to public addresses using NAT, but your Mac OS X Server gateway might be routing information between various private address subnets. Likewise, you might want to run a firewall between network segments in your own LAN.

Any condition in which you'd want to route network traffic through the server without masquerading IP addresses is a condition that 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 network ports must be properly configured and the Firewall service must be enabled.

To configure a gateway without NAT service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select NAT.
4. Click Settings.
5. Select "IP Forwarding only."
6. Click Save.

Monitoring NAT Service

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

Viewing the NAT Status Overview

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

To see the overview:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select NAT.
4. Click Overview to see whether the service is running, when it started, and the number of TCP, UDP, and ICMP links.

Common Network Administration Tasks That Use NAT

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

Linking a LAN to the Internet Through One IP Address

To link a LAN, you need a Mac OS X Server computer with two network interfaces: one to connect to the Internet and one to connect to your private network. The steps below use the following configuration as an example:

- **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 is provided by your ISP)
- **Internet or public DNS IP address***:* 17.254.1.6 (example only, your IP number is provided by your ISP)
- **Private network IP address range and netmask***:* 192.168.0.2–192.168.0.254 (also expressed as192.168.0.0/24 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

  This last setting is not required because NAT can be used with static IP addresses instead of DHCP. However, configuring this setting makes it easier to configure computers.

To configure your NAT LAN:

1. On the gateway server, open the Network pane of System Preferences.
2. In the active Network screen, make sure the interface "Built-in Ethernet" is at the top of the list of interfaces; if not, drag it to the top of the list.

   This sets the default gateway in the routing table. The top interface is always configured for 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 are:

   - 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 are:

   - IP address: 192.168.0.1
   - Netmask: 255.255.255.0
   - DNS: 17.254.1.6
5. If necessary, click Apply Now.
6. Open Server Admin and connect to the server.
7. Click the triangle to the left of the server.

   The list of services appears.
8. From the expanded Servers list, select DHCP.
9. Click Subnets and create a subnet 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: 17.254.1.6

   For detailed information about configuring DHCP, see "Creating Subnets" on page 26.
10. To start DHCP service, click the Start DHCP button below the Servers list.
11. In Server Admin, choose NAT from the expanded Servers list.

12 Configure NAT using the following setting:

External network interface: en0

13 If necessary, click Save.

14 To start NAT service, click the Start NAT button below the Servers list.

15 In Server Admin, choose Firewall from the expanded Servers list.

16 Create firewall rules to permit access to and from your private network.

For example, create an IP address group named "Private LAN" for the addresses 192.168.0.0/16.

For more information, see "Creating an Address Group" on page 90.

17 To start Firewall service, click the Start Firewall button below the Servers list.

18 Start any services you want the private LAN to access (web, SSH, file sharing, and so on) using the "Private LAN" group.

For more information, see "Configuring Services Settings" on page 88.

19 Start any services you want the Internet to access on your private LAN (web, SSH, file sharing, and so on) using the "any" address group.

For more information, see "Configuring Services Settings" on page 88.

20 Click Save.

Setting Up a LAN Party for Gaming

Some Internet-enabled games allow multiple players to connect online over a LAN. This is known as a *LAN party*. Setting up a LAN party is essentially the same as the process found in "Linking a LAN to the Internet Through One IP Address" on page 118.

Special considerations:

- Open only the ports necessary to play an Internet-enabled game.
- If the game is played only inside the LAN, don't open the firewall to game ports.
- If you have computers joining and leaving the LAN, use DHCP 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, suppose you have a NAT gateway called domain.example.com with an address of 17.100.0.1 that is set to forward web traffic (port 80) to 10.0.0.5 (port 80) behind the firewall and that sends packet requests for ssh traffic (port 22) to 10.0.0.15 (port 22).

In this example, the NAT gateway is not really serving the web content (the server at 10.0.0.5 is) but the server at 10.0.0.5 is invisible to the clients browsing the web site.

Viewed from the Internet you have one server, but viewed from behind the NAT barrier, you have as many or as few as you need. You can use this setup for load balancing or as an organizational scheme for the network's topography.

Virtual servers also enable you to easily reroute network traffic to other computers on the LAN by reconfiguring the gateway.

Virtual servers require three service configurations:

- **NAT**: The NAT service must be configured with port forwarding of the desired virtual port.
- **DNS**: The DNS record for the server should accept a few aliases of common services and resolve them all to the same IP address.
- **Firewall**: The firewall must permit traffic on specific ports to have access to the NAT LAN.

In this example, you set up a NAT gateway and route 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

  This last setting in not required because NAT can be used with static IP addresses instead of DHCP. However, configuring this setting makes it easier to configure computers.

To configure virtual servers:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select DHCP.
4. Click Subnets and 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 192.168.0.3

For more information, see "Creating Subnets" on page 26 and "Assigning Static IP Addresses Using DHCP" on page 35.

5. To start DHCP service, click the Start DHCP button (below the Servers list).
6. In Server Admin, choose NAT from the expanded Servers list.
7. Configure NAT using the following setting:
   - **External network interface**: en0
   - **Port forwarding**: TCP port 80 (web) to 192.168.0.2
   - **Port forwarding**: TCP port 25 (mail) to 192.168.0.3
8. Click Save.
9. To start NAT Service, click the Start NAT button below the Servers list.
10. In Server Admin, choose Firewall from the expanded Servers list.
11. Create Firewall rules to permit access to your private network.

    For more information, see "Creating an Address Group" on page 90.
12. Enable the two services you want the Internet to access on your private LAN (web and SMTP mail) using the "any" address group.

    For more information, see "Configuring Services Settings" on page 88.
13. Click Save.
14. To start Firewall service, click the Start Firewall button (below the Servers list).
15. Contact your DNS provider (usually your ISP) to add two aliases to your gateway server's DNS record.

    Request an A record with the name www.example.com to the IP address 17.100.0.1.

    Request an MX record with the name mail.example.com to the same IP address.

    These records are in addition to existing A and CNAME records for your domain.

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

If you want to change the servers behind the NAT (for example, to perform a hardware upgrade), now all you need to do is change the DHCP static IP address to the Ethernet addresses of the new servers. The new servers are assigned the existing internal IP addresses designated for web and mail, and the gateway forwards the traffic to the new servers seamlessly.

Where to Find More Information

For More Information About natd

The daemon process that controls NAT service is `natd`. For information about how to access `natd` features and implement them, see the `natd` man page.

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 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 NAT descriptions, see:

- RFC 1631
- RFC 3022

Working with VPN Service

6

This chapter describes how to set up and manage VPN service in Mac OS X Server.

By configuring a Virtual Private Network (VPN) on your server you can give users a more secure way of remotely communicating with computers on your network.

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

A VPN consists of 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 securely connect users working away from the office (for example, at home) to the LAN through a connection such as the Internet. From the user's perspective, the VPN connection appears as a dedicated private link.

VPN technology can also connect an organization to 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 can permit verified mobile users to access private computer resources (file servers and so on) using any connection to the Internet. VPNs can also link multiple LANs together over great distances using the existing Internet infrastructure.

VPN and Security

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

Transport Protocols

There are two encrypted transport protocols: Layer Two Tunneling Protocol, Secure Internet Protocol (L2TP/IPSec) and Point–to–Point Tunneling Protocol (PPTP). You can enable either or both of these protocols. Each has its own strengths and requirements.

L2TP/IPSec

L2TP/IPSec uses strong IPSec encryption to tunnel data to and from network nodes. It is based on Cisco's L2F protocol.

IPSec requires security certificates (either self-signed or signed by a certificate authority such as Verisign) or a predefined shared secret between connecting nodes.

The shared secret must be entered on the server and the client.

The shared secret is not a password for authentication, nor does it generate encryption keys to establish secure tunnels between nodes. It is a token that the key management systems use to trust each other.

L2TP is Mac OS X Server's preferred VPN protocol because it has superior transport encryption and can be authenticated using Kerberos.

PPTP

PPTP is a commonly used Windows standard VPN protocol. PPTP offers good encryption (if strong passwords are used) and supports a number of authentication schemes. It uses the user-provided password to produce an encryption key.

By default, PPTP supports 128-bit (strong) encryption. PPTP also supports the 40-bit (weak) security encryption.

PPTP is necessary if you have Windows clients with versions earlier than Windows XP or if you have Mac OS X v10.2.x clients or earlier.

Authentication Method

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

Kerberos is a secure authentication protocol that uses a Kerberos Key Distribution Server as a trusted third party to authenticate a client to a server.

MS-CHAPv2 authentication encodes passwords when they're sent over the network, and stores them in a scrambled form on the server. This method offers good security during network transmission. It is also the standard Windows authentication scheme for VPN.

Mac OS X Server PPTP VPN can also use other authentication methods. Each method has its own strengths and requirements. These other authentication methods for PPTP are not available in Server Admin.

If you want to use an alternative authentication scheme (for example, to use RSA Security's SecurID authentication), you must edit the VPN configuration file manually. The configuration file is located at /Library/Preferences/SystemConfiguration/com.apple.RemoteAccessServers.plist

For more information, see "Offering SecurID Authentication with VPN Server" on page 138.

Using VPN Service with Users in a Third-Party LDAP Domain

To use VPN service for users in a third-party LDAP domain (an Active Directory or Linux OpenLDAP domain), you must be able to use Kerberos authentication. If you need to use MSCHAPv2 to authenticate users, you can't offer VPN service for users in a third-party LDAP domain.

Before You Set Up VPN Service

Before setting up VPN service, 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 v10.5 and v10.4.x clients	X	X
Mac OS X v10.3.x clients	X	X
Mac OS X v10.2.x clients		X
Windows clients	X (if Windows XP)	X
Linux or Unix clients	X	X

If you're using L2TP, you must have a Security Certificate (from a certificate authority or self-signed), or a predefined shared secret between connecting nodes. If you use a shared secret, it must also be secure (at least 8 alphanumeric characters, including punctuation and without spaces; preferably 12 or more) and kept secret by users.

If you're using PPTP, make sure all your clients support 128-bit PPTP connections for greatest transport security. Using only 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 of 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, "Working with DHCP Service," on page 25.

Enabling VPN also requires Firewall services to be configured. The firewall settings must be able to pass network traffic from external IP addresses through the firewall to the LAN. The firewall settings can be as open or restricted as necessary.

For example, if your VPN clients use a large range of IP addresses (you have many users, each connecting from different ISPs) you might 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. You must also open the relevant firewall ports for the VPN type you are using (L2TP or PPTP).

Further, a VPN using L2TP must permit traffic for VPN clients on UDP port 4500 (IKE NAT Traversal) if you are using a NAT gateway.

Your specific network configuration can also require other open ports.

Setup Overview

Here is an overview of the steps for setting up print service:

Step 1: Before you begin
For information to keep in mind before you setup VPN service, read "Before You Set Up VPN Service" on page 127 and "Configuring Other Network Services for VPN" on page 128.

Step 2: Turn VPN service on
Before configuring VPN service, you must turn it on. See "Turning VPN Service On" on page 129.

Step 3: Configure VPN L2TP settings
Use Server Admin to enable L2TP over IPSec, set the IP address allocation range, and set the shared secret or security certificate. See "Configuring L2TP Settings" on page 129.

Step 4: Configure VPN PPTP settings
Use Server Admin to enable PPTP to specify, encryption key length, and to specify the IP address allocation range. See "Configuring PPTP Settings" on page 130.

Step 5: Configure VPN Logging settings
Use the Logging settings to enable VPN verbose logging. See "Configuring Logging Settings" on page 132.

Step 6: Configure VPN Client Information settings
Use Server Admin to configure network settings for VPN clients. See "Configuring Client Information Settings" on page 132.

Turning VPN Service On

Before you can configure VPN service, you must turn the VPN service on in Server Admin.

To turn VPN service on:

1. Open Server Admin and connect to the server.
2. Click Settings, then click Services.
3. Click the VPN checkbox.
4. Click Save.

Setting Up VPN Service

There are two groups of settings for VPN service in Server Admin:

- **Connections.** Shows you information about users who are connected using VPN.
- **Settings.** Configures and manages L2TP and PPTP VPN service connections.

The following sections describe how to configure these settings. A final section explains how to start VPN service after you set up VPN.

Configuring L2TP Settings

Use Server Admin to designate L2TP as the transport protocol.

If you enable 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 the group that will use the VPN service (if needed).

If both L2TP and PPTP are used, each protocol should have a separate, nonoverlapping address range.

When configuring VPN, make sure the firewall allows VPN traffic on needed ports with the following settings:

- For the "any" address group, enable GRE, ESP, VPN L2TP (port 1701), and VPN ISAKMP/IKE (port 500).
- For the "192.168-net" address group, choose to allow all traffic.

For more information, see "Configuring Services Settings" on page 88.

To configure L2TP settings:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click Settings, then click L2TP.
5. Select the "Enable L2TP over IPSec" checkbox.
6. In the "Starting IP address" field 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. In the "Ending IP address" field 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. (Optional) You can load-balance VPN by selecting the Enable Load Balancing checkbox and entering an IP address in the Cluster IP address field.
9. Choose a PPP authentication type.

   If you choose Directory Service and your computer is bound to a Kerberos authentication server, from the Authentication pop-up menu select Kerberos. Otherwise, choose MS-CHAPv2.

   If you choose RADIUS, enter the following information:

   Primary IP Address: Enter the IP address of the primary RADIUS server.

   Shared Secret: Enter a shared secret for the primary RADIUS server.

   Secondary IP Address: Enter the IP address of the secondary RADIUS server.

   Shared Secret: Enter a shared secret for the secondary RADIUS server.
10. Enter the shared secret or select the certificate to use in the IPSec Authentication section.

    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.
11. Click Save.

Configuring PPTP Settings

Use Server Admin to designate PPTP as the transport protocol.

If you enable this protocol, you must also configure connection settings. You should designate an encryption key length (40-bit or 128-bit), the IP address allocation range to be given to your clients, and the group that will use the VPN service (if needed).

If you use L2TP and PPTP, each protocol should have a separate, nonoverlapping address range.

When configuring VPN, make sure the firewall allows VPN traffic on needed ports with the following settings:

- For the "any" address group, enable GRE, ESP, VPN L2TP (port 1701), and IKE (port 500).
- For the "192.168-net" address group, choose to allow all traffic.

For more information, see "Configuring Services Settings" on page 88.

To configure PPTP settings:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click Settings, then click PPTP.
5. Select "Enable PPTP."
6. If needed, select "Allow 40-bit encryption keys in addition to 128-bit" to permit both 40-bit and 128-bit key encryption access to VPN.

> ***WARNING:*** 40-bit encryption keys are much less secure but can be necessary for some VPN client applications.

7. In the "Starting IP address" field set the beginning IP address of the VPN allocation range.

   It can't overlap the DHCP allocation range, so enter 192.168.0.128.
8. In the "Ending IP address" field set the ending IP address of the VPN allocation range.

   It can't overlap the DHCP allocation range, so enter 192.168.0.255.
9. Choose a PPP authentication type.

   If you choose Directory Service and your computer is bound to a Kerberos authentication server, from the Authentication pop-up menu select Kerberos. Otherwise, choose MS-CHAPv2.

   If you choose RADIUS, enter the following information:

   Primary IP Address: Enter the IP address of the primary RADIUS server.

Shared Secret: Enter a shared secret for the primary RADIUS server.

Secondary IP Address: Enter the IP address of the secondary RADIUS server.

Shared Secret: Enter a shared secret for the secondary RADIUS server.

10 Click Save.

Configuring Client Information Settings

When a user connects 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 must configure the network mask, DNS address, and search domains.

To configure Client Information settings:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of servers appears.

3 From the expanded Servers list, select VPN.

4 Click Settings, then click Client Information.

5 Enter the IP address of the DNS server.

Add the gateway computer's internal IP address (usually something like 192.168.x.1).

6 Enter search domains as needed.

7 Click Save.

Configuring Logging Settings

You can choose from two levels of detail for VPN service logs.

- **Nonverbose logs**: Describe conditions where you must take immediate action (for example, if the VPN service can't start up).
- **Verbose logs**: Record all activity by the VPN service, including routine functions.

By default, nonverbose logging is enabled.

To change logging settings to verbose:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of servers appears.

3 From the expanded Servers list, select VPN.

4 Click Settings, then click Logging.

5 Select "Verbose logging" to enable verbose logging.

6 Click Save.

Starting VPN Service

You use Server Admin to start VPN service.

To start VPN service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click the Start VPN button below the Servers list.

   Click Settings, then click L2TP or PPTP and verify that the "Enable L2TP over IPsec" or "Enable PPTP" checkbox is selected.

Managing VPN Service

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

Stopping VPN Service

You use Server Admin to stop VPN service.

To stop VPN service:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click the Stop VPN button below the Servers list.

Configuring VPN Network Routing Definitions

By using network routing definitions, you can choose whether to route data from VPN clients to an address group through the VPN tunnel (referred to as *private*) or over the VPN user's ISP connection (referred to as *public*).

For example, you can have all VPN client traffic that goes to the LAN IP address range go through the secure tunnel to the LAN, but make all traffic to other addresses be routed through the user's normal, unsecured Internet connection.

This helps you have greater control over what goes through the VPN tunnel.

Important Notes About VPN Routing Definitions

- If no routing definitions are added, traffic is routed through the VPN connection by default.

- If routing definitions are added, the VPN connection is no longer set as the default route, and traffic destined for addresses not specifically declared as a private route will not go over the VPN connection.
- DNS lookups go over the VPN connection regardless of the routes that are set.
- Definitions are unordered. They only apply the description that most closely matches the packet being routed.

Example

Suppose your LAN's IP addresses are 17.x.x.x addresses. If you make no routing definitions, every VPN client's network traffic (such as web browser URL requests, LPR printer queue print jobs, and file server browsing) is routed from the client computer through the VPN tunnel to the 17.x.x.x LAN.

You decide that you don't want to manage all 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 goes through the client computer's normal Internet connection.

To limit the traffic the VPN tunnel handles, enter a routing definition designating traffic to the 17.x.x.x network as private, which sends it through the VPN tunnel. In the routing definition table you'd enter 17.0.0.0 255.0.0.0 Private.

All traffic to the LAN is now sent over the VPN connection and, by default, all other addresses not in the definitions table are sent over the client computer's unencrypted Internet connection.

You then decide that there are 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 through the client computer's Internet connection and not pass through the VPN tunnel. The addresses might be outside the firewall and not accessible from the 17.x.x.x LAN.

As an example, if you want to use addresses in the range 17.100.100.x, you would enter an extra routing definition as follows: 17.100.100.0 255.255.255.0 Public.

Because 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 through the client computer's Internet connection.

In summary, if you add routes, any routes you specify as private go over the VPN connection, and any declared as public do not go over the VPN connection. All others not specified also do not go over the VPN connection.

To set routing definitions:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.

3 From the expanded Servers list, select VPN.

4 Click Settings, then click Client Information.

5 Click the Add (+) button.

6 Enter a destination address range of the packets to be routed by specifying:

A base address (for example, 192.168.0.0)

A network mask (for example, 255.255.0.0)

7 From the Type pop-up menu, select the routing destination.

Private means to route client traffic through the VPN tunnel.

Public means to use the normal interface with no tunnel.

8 Click OK.

9 Click Save.

Limiting VPN Access to Specific Users or Groups

By default, all users on the server or in the master directory have access to the VPN when it is enabled. You can limit VPN access to specific users for security or ease of administration. You can limit access to VPN by using Mac OS X Server's Access Control List (ACL) feature.

ACLs allow you to designate service access to users or groups on an individual basis. For example, you can use an ACL to permit a user to access a specific file server or shell login, while denying access to all other users on the server.

To limit VPN access using ACLs:

1 Open Server Admin and connect to the server.

2 Click Settings, then click Access.

3 Click Services.

4 Select "For selected services below."

5 In the service access list, select VPN.

6 Select "Allow only users and group below."

7 To reveal a Users and Groups drawer, click the Add (+) button.

8 Drag users or groups to the access list.

9 Click Save.

Limiting VPN Access to Specific Incoming IP Addresses

By default, Firewall service blocks incoming VPN connections, but you can provide limited VPN access to certain IP addresses for security or ease of administration.

You can limit access to the VPN by using Mac OS X Server's Firewall service. When configuring the firewall for L2TP and PPTP you must configure GRE, ESP, and IKE to permit VPN access through the firewall.

To limit VPN access by IP address:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select Firewall.
4. Click Settings.
5. Select Advanced, then click the Add (+) button.
6. From the Action pop-up menu, choose "Allow."
7. From the Protocol pop-up menu, choose an option.

   If you use L2TP for VPN access, choose UDP.

   If you use PPTP for VPN access, choose TCP.
8. From the Service pop-up menu, choose VPN L2TP or VPN PPTP.

   The relevant destination port is added to the Port field.
9. (Optional) Select the "Log all packets matching this rule" checkbox.
10. From the address pop-up menu of the Source section, choose Other and enter the source IP address range (using CIDR notation) that you want to give access to VPN.

    You can also specify a port in the Port field of the Source section.

    Computers that have an IP address in the IP address range that you specified in the source IP address field, communicating on the source port you specified, can connect to the VPN service.
11. From the Destination Address pop-up menu, choose the address group that contains 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).
12. From the Interface pop-up menu that this rule applies to, choose "In."

    "In" refers to the packets coming into the server.
13. Click OK.
14. Click the Add (+) button.
15. From the Action pop-up menu, choose "Allow."
16. From the Protocol pop-up menu, choose a protocol or Other.

    If you are adding GRE or ESP, choose Other and enter "any" in the field.

If you are adding VPN ISAKMP/IKE, choose UDP.

17 From the Service pop-up menu, choose a service.

If you are adding GRE, choose "GRE - Generic Routing Encapsulation protocol."

If you are adding ESP, choose "ESP - Encapsulating Security Payload protocol."

If you are adding VPN ISAKMP/IKE, choose "VPN ISAKMP/IKE." Destination port 500 is added to the Port field.

18 From the Address pop-up menu of the Source section, choose "any."

19 In the Port field of the Source section, enter "any."

20 From the Address pop-up menu of the Destination section, choose "any."

21 In the Port field of the Destination section, enter a port number.

If you are adding VPN ISAKMP/IKE, enter 500 if it is not already shown.

22 From the Interface pop-up menu, choose "Other" and enter "any" in the Other field of the Interface section.

23 Click OK.

24 Repeat steps 14 through 23 for GRE, ESP, and VPN ISAKMP/IKE.

25 Click Save to apply the filter immediately.

Supplementary Configuration Instructions

The following section describes procedures for optional scenarios. They require integration with an existing directory service or with third-party authentication services.

Enabling VPN-PPTP Access for Users in an LDAP Domain

In Mac OS X v10.5, you can use a command-line tool to enable PPTP-VPN connections for users in an LDAP domain.

This resolves a situation where users can establish a VPN connection using PPTP to a Mac OS X Server that, when established, is not used by network traffic. This situation affects Mac OS X Server v10.3, v10.4, and v10.5.

To enable VPN-PPTP access for users in an LDAP domain:

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

For example, if the server running the VPN service is the LDAP master, enter the following command in Terminal:

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

If the server 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 address is 17.221.67.87, enter the following command in Terminal:

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

2 When prompted, enter the username and password.

If the VPN server is the LDAP master, enter the administrator name and password of the server.

If the LDAP directory is on a different server, enter the administrator name and password of the server that hosts the LDAP directory (or the administrator name and password used to add users to the LDAP directory in Workgroup Manager).

The tool adds a user to the LDAP directory and sets up configuration elements in the VPN Server so it can support PPTP.

3 In the VPN Service Settings pane of Server Admin, configure PPTP.

4 Start VPN Service.

Offering SecurID Authentication with VPN Server

RSA Security provides strong authentication. It uses hardware and software tokens to verify user identity. SecurID authentication is available for L2TP and PPTP transports. For details and product offerings, see www.rsasecurity.com.

VPN service supports SecurID authentication but it cannot be set up from Server Admin. If you choose this authentication tool, you must change the VPN configuration manually.

Set up SecurID:

1 From your SecurID server, copy the sdconf.rec file to a new folder on your Mac OS X Server named /var/ace.

There are several ways to do this. The following illustrates one method:

a Open Terminal (/Applications/Utilities/).

b Enter `sudo mkdir /var/ace`.

c Enter your administrator password.

d In the Dock, click Finder.

e From the Go menu, choose Go > Go to Folder.

f Enter: `/var/ace`.

g Click Go.

h From your SecurID server, copy the sdconf.rec file into the "ace" folder.

i If you see a dialog indicating that the "ace" folder cannot be modified, click Authenticate to permit the copy.

2 Enable EAP-SecurID authentication on your VPN service for the protocols you want to use it with.

To use it with PPTP, enter these two commands in Terminal (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, enter these two commands in Terminal (each only one line):

```
# 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"
```

3 Complete the remaining VPN service configuration tasks using Server Admin.

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 for enabled VPN services. It tells you how many L2TP and PPTP clients are connected, which authentication method is selected, and when the service was started.

To view the overview:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of servers appears.

3 From the expanded Servers list, select VPN.

4 Click Overview.

Changing the Log Detail Level for VPN Service

You can choose from two levels of detail for VPN service logs:

- **Nonverbose**: These logs describe only conditions where you must take immediate action (for example, if the VPN service can't start up).
- **Verbose**: These logs record all activity by the VPN service, including routine functions.

By default nonverbose logging is enabled.

To change the VPN log detail to verbose

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click Settings, then click Logging.
5. Select "Verbose logging" to enable verbose logging.
6. Click Save.

Viewing the VPN Log

Monitoring VPN logs helps you make sure your VPN is running properly. VPN logs can help you troubleshoot problems. The log view shows the contents of the /var/log/ppp/vpnd.log file. You can filter the log records with the text filter box in the Log pane of VPN.

To view the log:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click Log.

Viewing VPN Client Connections

You can monitor VPN client connections to maintain secure access to the VPN. By viewing the client connection screen, you can see:

- Users connected
- IP address users are connecting from
- IP address your network assigned to users
- Type and duration of connections

You can sort the list by clicking the column headers.

To view client connections:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click Connections.

Common Network Administration Tasks That Use VPN

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

Linking a Computer at Home with a Remote Network

You can use VPN to link a computer to a remote network, giving you access to it as if it were physically connected to the LAN. The following is an example of a linked computer configuration:

- **User authentication**: 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 as192.168.0.0/24 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 that can connect to a remote LAN using L2TP, with full access rights.

Step 1: Configure VPN

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of servers appears.
3. From the expanded Servers list, select VPN.
4. Click Settings, then click L2TP.
5. Enable L2TP over IPsec.
6. In the "Starting IP address" field 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. In the "Ending IP address" field 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. In the IPSec Authentication section 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.
9. Click Save.
10. Click Client Information.

11 Enter the IP address of the internal LAN DNS server (192.168.0.2).

12 Leave routing definitions empty.

All traffic from the client will go through the VPN tunnel.

13 Click Save.

14 Click Start VPN below the Servers list.

Step 2: Configure the firewall

1 Create an address group for the VPN allocation range.

For more information, see "Creating an Address Group" on page 90.

2 Open the firewall to external VPN connections by enabling L2TP connections in the "any" address group.

For more information, see "Configuring Services Settings" on page 88.

3 Configure the firewall for the VPN address group, permitting or denying ports and services as needed.

4 Save your changes.

5 Start or restart the firewall.

Step 3: Configure the client

This example is of a Mac OS X client using Network preferences.

1 Open System Preferences, then click Network.

2 Click the Add (+) button at the bottom of the network connection services list and then choose VPN from the Interface pop-up menu.

3 From the VPN Type pop-up menu, choose "L2TP over IPSec".

4 Enter a VPN service name in the Service Name field, then click Create.

5 Enter the DNS name or IP address in the Server Address field.

Server Address: gateway.example.com

Account Name: <*the user's short name*>

6 Click Authentication Settings and enter the following configuration information:

User Authentication: Use Password <*user's password*>

Machine Authentication: Use Shared Secret <prDwkj49fd!254>

7 Click OK.

The user can now connect.

Accessing a Computing Asset Behind a Remote Network Firewall

Accessing a single computing asset behind a firewall differs from permitting a client computer 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 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 information in the section "Linking a Computer at Home with a Remote Network" on page 141, and adds:

- **File server IP address**: 192.168.0.15
- **File server type**: Apple File Sharing

For this scenario, the procedure is similar to that use for "Linking a Computer at Home with a Remote Network" on page 141, 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 who are now logged in through the VPN gateway can access the file server, and no other network traffic can go through the encrypted gateway.

Linking Two or More Remote Network Sites

You can use a VPN to link a computer to a main network, and you can also link networks.

When two networks are linked they can interact as if they are physically connected. Each site must have its own connection to the Internet but the private data is sent encrypted between the sites.

This type of link is useful for connecting satellite offices to an 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 named `s2svpnadmin` ("site-to-site VPN admin").

Using `s2svpnadmin` requires the use of (and facility with) the Terminal, and the administrator must have access to root privileges through `sudo`. For more about `s2svpnadmin`, see the `s2svpnadmin` man page.

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

Site-to-site VPN connections can be only made using L2TP/IPSec VPN connections. You cannot link two sites using PPTP and these instructions.

This example uses the following settings:

- Desired VPN type: L2TP
- Authentication: Using 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 as192.168.0.0/16 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 as192.168.20.0/24 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 using L2TP.

Step 1: Run s2svpnadmin on both site gateways

1. Open Terminal and start `s2svpnadmin` by entering:

   ```
   $ sudo s2svpnadmin
   ```

2. Enter the relevant number for "Configure a new site-to-site server."
3. Enter an identifying configuration name (no spaces permitted).

   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 that you want to use.
7. Enter at least one addressing policy for the configuration.
8. Enter a local subnet network address (for example, 192.168.0.0 for site 1, 192.168.20.0 for site 2).
9. For the address range, enter the prefix bits in CIDR notation.

In this example, the CIDR notation for the subnet range is 192.168.2.0/24 for site 1, so you would enter 24.

10 Enter a remote subnet network address (for example, 192.168.20.0 for site 1, 192.168.0.0 for site 2).

11 For the address range, enter the prefix bits in CIDR notation.

In this example, the CIDR notation for the subnet range is 192.168.2.0/24 for site 1, so you would enter 24.

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.

Repeat policy steps 7 through 12 for the new policies.

13 Press "y" to enable the site configuration.

You can verify 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 gateways

1 Create an address group for each server with only the server's public IP address.

In this example, name the first group Site 1 and enter the public IP address of the server. Then name the second group Site 2 and enter the public IP address of the other server.

For more information, see "Creating an Address Group" on page 90.

2 Open the firewall to external VPN connections by enabling L2TP (port 1701) connections and IKE NAT Traversal (port 4500) in the "any" address group.

For more information, see "Configuring Services Settings" on page 88.

3 Create the following Advanced IP filter rules on both site gateways:

Filter Rule 1	Setting
Action:	Allow
Protocol:	UDP
Source Address:	Site 1
Destination Address:	Site 2
Interface:	Other, enter "isakmp"

Filter Rule 2	Setting
Action:	Allow
Protocol:	UDP

Filter Rule 2	Setting
Source Address:	Site 2
Destination Address:	Site 1
Interface:	Other, enter “isakmp”

Filter Rule 3	Setting
Action:	Allow
Protocol:	Other, enter “esp”
Source Address:	Site 1
Destination Address:	Site 2

Filter Rule 4	Setting
Action:	Allow
Protocol:	Other, enter “esp”
Source Address:	Site 2
Destination Address:	Site 1

Filter Rule 5	Setting
Action:	Allow
Protocol:	Other, enter “ipencap”
Source Address:	Site 1
Destination Address:	Site 2

Filter Rule 6	Setting
Action:	Allow
Protocol:	Other, enter “ipencap”
Source Address:	Site 2
Destination Address:	Site 1

Filter Rule 7	Setting
Action:	Allow
Protocol:	Other, enter “gre”
Source Address:	Site 1
Destination Address:	Site 2

Filter Rule 8	Setting
Action:	Allow
Protocol:	Other, enter “gre”

Filter Rule 8	Setting
Source Address:	Site 2
Destination Address:	Site 1

For more information about creating advanced rules, see "Configuring Advanced Firewall Rules" on page 93.

These rules permit the encrypted traffic to be passed to both hosts.

4 Save your changes.

5 Start or restart the firewall, as needed.

Step 3: Start VPN service on both site gateways

1 For both VPN gateways, open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 Select VPN from the expanded Servers list.

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

4 Click Start VPN.

You should now be able to access a computer on the remote LAN from the local LAN. To verify the link, use `ping` or some other means.

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 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 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.

Working with RADIUS Service

7

By configuring a RADIUS server with Open Directory you can secure your wireless environment from unauthorized users.

Wireless networking gives companies greater network flexibility, seamlessly connecting laptop users to the network and giving them the freedom to move within the company while staying connected to the network.

This chapter describes how to configure and use Remote Authentication Dial In User Service (RADIUS) to keep your wireless network secure and to make sure it is used only by authorized users.

RADIUS is used to authorize Open Directory users and groups so they can access AirPort Base Stations on a network. By configuring RADIUS and Open Directory you can control who has access to your wireless network.

RADIUS works with Open Directory and Password Server to grant authorized users access to the network through an AirPort Base Station. When a user attempts to access an AirPort Base Station, AirPort communicates with the RADIUS server using the Extensible Authentication Protocol (EAP) to authenticate and authorize the user.

Users are given access to the network if their user credentials are valid and they are authorized to use the AirPort Base Station. If a user is not authorized, he or she cannot access the network through the AirPort Base Station.

Before You Set Up RADIUS Service

This section contains information you should consider before setting up RADIUS on your network.

Setting Up RADIUS Service for the First Time

If you're setting up your own RADIUS server, follow the steps in this section.

Step 1: Turn RADIUS service on
Before configuring service, turn on RADIUS. See "Turning RADIUS Service On" on page 150.

Step 2: Add AirPort Base Stations to a RADIUS server
Decide which AirPort Base Stations you want to add to the RADIUS server. See "Adding AirPort Base Stations to a RADIUS Server" on page 152.

Step 3: Remotely configure an AirPort Base Station
Use Server Admin to configure AirPort Base Stations. See "Remotely Configuring AirPort Base Stations" on page 152.

Step 4: Configure RADIUS to use certificates
Use Server Admin to configure RADIUS to use certificates to trust Base Stations. See "Remotely Configuring AirPort Base Stations" on page 152.

Step 5: Start RADIUS service
To start the RADIUS service, see "Starting or Stopping RADIUS Service" on page 153.

Turning RADIUS Service On

Before you can configure RADIUS settings, you must turn on RADIUS service in Server Admin.

To turn RADIUS service on:

1. Open Server Admin and connect to the server.
2. Click Settings, then click Services.
3. Select the RADIUS checkbox.
4. Click Save.

Setting Up RADIUS Service

This section describes how to add AirPort Base Stations to your RADIUS server, configure AirPort Base Stations remotely, and configure RADIUS to use certificates to trust AirPort Base Stations.

The following sections describe how to configure these settings. A final section tells you how to start the RADIUS service when you finish.

Configuring RADIUS Using the Configuration Assistant

Mac OS X Server v10.5 offers a configuration assistant for the RADIUS service. The configuration assistant guides you thought the RADIUS configuration process and start RADIUS.

To configure RADIUS using the configuration assistant:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select RADIUS.
4. Click Overview.
5. Click Configure RADIUS Service.
6. In the RADIUS Server Certificate pane, select one of the following:

   If you select "Choose an existing certificate," choose the certificate you want to use from the pop-up menu and click Continue.

   If you want to create a self-signed certificate, select "Create a self-signed certificate," and click Continue.

   a. Fill out identity information.

      The common name is the fully qualified domain name of the server that uses SSL enabled services.

   b. Enter starting and ending validity dates.

   c. Select a private key size (1024 bits is the default).

   d. Enter a passphrase for the private key and Click Save.

      This passphrase should be more secure than a normal password. You should use at least 20 characters, including mixed case, numbers, and punctuation. Don't repeat characters and don't use words contained in the dictionary.

   e. Click Continue.

      A message appears explaining self-signed certificates.

   f. Click Continue.
7. From the Available Base Stations list, select the Base Station you want and click Add.

   If the Base Station has a password enter it in the Base Station Password field, then click Add.

   If you want to remove a Base Station from the Selected Base Stations list, select it and click Remove.
8. Click Continue.
9. In the RADIUS Allow Users pane you can restrict user access.

   If you select the "Allow all users," all users will have access to the Base Stations you selected.

   If you select "Restrict to members of group," only users of a group can access the Base Stations you selected.
10. Click Continue.

11 In the RADIUS setting confirmation pane, verify your settings are correct.

You can also print or save you RADIUS configuration settings.

12 Click Confirm.

Adding AirPort Base Stations to a RADIUS Server

You use the Settings pane of RADIUS in Server Admin to add AirPort Base Stations that will use the RADIUS service. You can add up to 64 Base Stations to RADIUS.

To add AirPort Base Stations to a RADIUS server:

1 On the management computer, open Server Admin.

2 Click the triangle to the left of the server.

The list of services appears.

3 In the expanded Servers list, click RADIUS.

4 Click Base Stations.

5 Below the AirPort Base Stations list, click Add.

6 Enter the following AirPort Base Station information:

Name: Specify the name of the AirPort Base Station.

Type: Specify the model of the AirPort Base Station.

IP Address*:* Specify the IP address of the AirPort Base Station.

Shared Secret and **Verify**: The shared secret is not a password for authentication, nor does it generate encryption keys to establish secure tunnels between nodes. It is a token that the key management systems uses to trust each other. The shared secret must be entered on the server as well as a client.

7 Click Add.

Remotely Configuring AirPort Base Stations

You can remotely configure AirPort Base Stations to use a RADIUS server in Server Admin.

To remotely configure AirPort Base Stations to use a RADIUS server:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select RADIUS.

4 Click Base Stations.

5 Highlight the AirPort Base Station in the AirPort Base Stations list, then click Edit.

If prompted for a password, enter the AirPort administrator password.

6 Click OK.

Configuring RADIUS to Use Certificates

You can use Server Admin to configure RADIUS to use custom certificates. Using a certificate increases the security and manageability of AirPort Base Stations.

To use a custom certificate:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select RADIUS.
4. Click Settings.
5. From the RADIUS Certificate pop-up menu, choose a certificate.

   If you have a custom certificate, choose Custom Configuration from the Certificate pop-up menu and enter the path to the certificate file, private key file, and certificate authority file. If the private key is encrypted, enter the private key passphrase and click OK.

   If you don't have a certificate and want to create one, click Manage Certificates. For more information about creating certificates, see *Server Administration*
6. Click Save.

Archiving RADIUS Service Logs

RADIUS service creates entries in the system log for error and alert messages. You can archive these log entries by using Server Admin.

To archive logs:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select RADIUS.
4. Click Settings.
5. Select the "Archive radiusd log for the past ___ days" checkbox and enter the number of days you want to archive.
6. Click Save.

Starting or Stopping RADIUS Service

You use Server Admin to start or stop RADIUS service. When stopping the service make sure no users are connected to AirPort Base Stations your RADIUS server manages.

To start or stop RADIUS service:

1. Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select RADIUS.

4 Click Start RADIUS or Stop RADIUS below the Servers list.

The service can take a few seconds to start or stop.

Managing RADIUS Service

This section describes tasks you might perform after you set up RADIUS service on your server.

Checking RADIUS Service Status

You can use Server Admin to check the status of RADIUS service.

To check RADIUS service status:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select RADIUS.

4 Click Overview to see whether the service is running, the number of client base stations, and when it was started.

Viewing RADIUS Service Logs

RADIUS service creates entries in the system log for error and alert messages. You can filter the log to narrow the number of viewable log entries and make it easier to find the entry you want to see.

To view logs:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select RADIUS.

4 Click Logs.

5 Choose a log to view (radiusconfig or radiusd).

To search for specific entries, use the Filter field below the log.

Editing RADIUS Access

You can restrict access to the RADIUS service by creating a group of users and adding them to the service access control list (SACL) of RADIUS.

To edit RADIUS access:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select RADIUS.
4. Click Settings, then click Edit Allowed Users.
5. Select "For selected services below," then select RADIUS.
6. Select "Allow only users and groups below."
7. Click the Add (+) button.
8. From the Users and Groups list, drag users or groups of users to the "Allow only users and groups below" list.

   If you want to remove users from the "Allow only users and groups below" list, select the users or groups of users and click the Delete (-) button.

   Only users in the list can use the RADIUS services.

Deleting AirPort Base Stations

You can use Server Admin to delete AirPort Base Stations from the RADIUS server.

When deleting AirPort Base Stations make sure the stations are disconnected from the network. Otherwise, unauthorized users might be able to access your network.

To delete AirPort Base Stations:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.
3. From the expanded Servers list, select RADIUS.
4. Click Base Stations.
5. In the AirPort Base Station list, highlight a Base Station and click Remove.
6. Verify you want to remove the Base Station by clicking Remove again.

Editing an AirPort Base Station Record

You can user Server Admin to edit an AirPort Base Station record on your RADIUS server.

To edit an AirPort Base Station record:

1. Open Server Admin and connect to the server.
2. Click the triangle to the left of the server.

   The list of services appears.

3 From the expanded Servers list, select RADIUS.

4 Click Base Stations.

5 In the AirPort Base Station list, highlight the Base Station you want to modify and click the Edit button.

6 Modify the Base Station information and click Save.

Saving an AirPort Base Station Internet Connect File

You can use Server Admin to save an AirPort Base Station internet connect file.

To save an AirPort Base Station Internet connect file:

1 Open Server Admin and connect to the server.

2 Click the triangle to the left of the server.

The list of services appears.

3 From the expanded Servers list, select RADIUS.

4 Click Base Stations.

5 In the AirPort Base Station list, highlight the base station.

6 Click Save Internet Connect File.

7 In the Save As field, enter the name.

8 From the Where pop-up menu, choose the location to save the file.

9 In the Wireless Network Name (SSID) field, enter the wireless network name.

10 Click Save.

Working with NTP Service

8

Using NTP service for time synchronization is important for reducing confusion that can be caused if time stamps are out of sync.

From shared file systems to billing services, correct timekeeping is a necessity. However, clocks on computers throughout a network can have widely different time stamps. Network Time Protocol (NTP) is used to synchronize the clocks in networked computers to a reference clock. NTP helps make sure that all computers on a network report the same time.

If an isolated network (or even a single computer) is unsynchronized, services that use time and date stamps (such as mail service, or web service with timed cookies) send wrong time and date stamps and are 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 referred to as "atomic clocks."

On the Internet, authoritative NTP servers (known as *Stratum 1* servers) keep track of the current UTC time. Other subordinate servers (known as *Stratum 2 and 3* servers) regularly query the Stratum 1 servers and estimate the time taken to send and receive the query. They then factor this estimate with the query result to set the Stratum 2 or 3 servers' time. The estimates are correct to the nanosecond.

Your LAN can then query Stratum 3 servers for the time. An NTP client computer on your network then takes the UTC time reference and converts it using 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 as an NTP client, receiving authoritative time from an Internet time server, and as an authoritative time server for a network. Your local clients can query your server to set their clocks.

If you set your server to answer time queries, set it to also query an authoritative time server on the Internet.

Setting Up NTP Service

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

Make sure your firewall permits NTP queries to an authoritative time server on UDP port 123, and incoming queries from local clients on the same port. For more information, see Chapter 4, "Working with Firewall Service."

To set up NTP service:

1. Open Server Admin and connect to the server.
2. Click Settings, then click Date & Time.
3. Make sure your server is configured to "Set date & time automatically."
4. From the pop-up menu, choose the server you want to act as a time server.
5. Click General.
6. Select the "Network Time Server (NTP)" checkbox.
7. Click Save.

Configuring NTP Service on Clients

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

Use the following instructions to set your clients to query your time server.

To configure NTP on clients:

1. Open System Preferences.
2. Click Date & Time.
3. Select the "Set date & time automatically" checkbox.
4. Select and delete the text in the field rather than using the pop-up menu.
5. Enter the host name of your time server.

Your host name can be either a domain name (such as time.example.com) or an IP address.

6 Close System Preferences.

Where to Find More Information

The working group, documentation, and FAQ for NTP can be found at www.ntp.org.

Listings of publicly accessible NTP servers and their use policies can be found at support.ntp.org/bin/view/Servers/WebHome.

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 technical details about a protocol in its RFC document.

You can search for RFC documents by number at www.ietf.org/rfc.html.

The official specification of NTP version 3 is RFC 1305.

Supporting a VLAN

9

Using a virtual local area network (VLAN) can prevent delays and data loss in environments with extremely high amounts of network traffic.

VLANs enable multiple computers on different physical LANs to communicate with each other as if they were on the same LAN.

Benefits include more efficient network bandwidth use and greater security, because broadcast or multicast traffic is only sent to computers on the common network segment.

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

Xserve G5 VLAN support conforms to the IEEE standard 802.1q.

Setting Up Client Membership for a VLAN

To set up and manage VLANs, you use the VLAN area of the Network pane of System Preferences.

Be sure that ports used by non-VLAN devices (non-802.1q-compliant) are configured to transmit untagged frames. If a noncompliant Ethernet device receives a tagged frame, it cannot understand the VLAN tag and drops the frame.

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

To set up a VLAN:

1. Log in to your server as an administrator.
2. Open the Network pane of System Preferences.
3. Click the Action pop-up menu and select Manage Virtual Interfaces.
4. Click the Add (+) button and select New VLAN.

5 In the VLAN Name field, enter a name for the VLAN.

6 In the Tag field, enter a tag (a number between 1 and 4094).

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

7 Select the Interface.

8 Click Create.

9 Click Done.

Where to Find More Information

For More Information About VLANs on the Internet
See www.ieee.org. The VLAN standard is defined by IEEE.

Reference Document
A reference document provides an overview of a protocol and includes 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.

The reference document is available at standards.ieee.org/getieee802/download/802.1Q-1998.pdf.

Supporting IPv6

10

Internet Protocol Version 6 (IPv6) is the Internet's next generation protocol, designed to replace the current Internet protocol.

The current Internet protocol, IP version 4 (IPv4, or just IP), is beginning to have problems coping with the growth and popularity of the Internet. The main problems for IPv4 are:

- **Limited IP addressing:** IPv4 addresses use 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 required to route IPv4 information is rapidly increasing as more computers connect to the Internet at an increasing rate.
- **End-to-end communication that's routinely circumvented:** This problem 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 two network end points. However, this frustrates a number of network services and is limiting.

IPv6 fixes some of these problems and helps prevents others. It improves routing and network autoconfiguration, it increases the number of network addresses to over $3x10^{38}$, and it 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 addressing:

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

A number of command-line tools installed with Mac OS X Server support IPv6 (for example, ping6 and traceroute6).

Support for IPv6 Addresses in Server Admin

The services above support IPv6 addresses, but not in Server Admin. IPv6 addresses fail if entered in IP address fields in Server Admin. IPv6 addresses for these services can be configured with command-line tools and by editing configuration files.

IPv6 Addresses

IPv6 addresses are different from IPv4 addresses. There are changes in address notation, reserved addresses, the address model, and address types.

Notation

IPv4 addresses are 4 bytes long and are expressed in decimals, but 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

The address is split into pairs of bytes, separated by colons. Each byte is represented as a pair of hexadecimal numbers. The following address is in IPv6 format:

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

This can be abbreviated as follows:

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

IPv6 addresses often contain 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

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 IPv4 notation. Using this mixed notation, the above example can be expressed as:

E3C5::4AC8:192.168.100.32

IPv6 Reserved Addresses

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

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

0:0:0:0:0:0:0:1 (loopback address, like127.0.0.1 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.

Routers don't need an IPv6 address, eliminating the need to configure routers for point-to-point unicasts.

IPv6 doesn't use IPv4 address classes.

IPv6 Address Types

IPv6 supports the following IP address types:

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

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. Multicast delivers messages to all nodes in the multicast group, but anycast delivers messages to one node in the multicast group.

Creating an IPv4 to IPv6 Gateway

Mac OS X Server includes an IPv4-to-IPv6 gateway that enables the deployment of IPv4-based server services in IPv6 networks to support the industry-wide IP transition.

You can configure the IPv4 to IPv6 gateway by setting up "6 to 4" in Network Preferences of your server.

Important: You must have a public IPv4 address (an IP address issued by your ISP) and you cannot be behind a gateway or NAT.

To configure a "6 to 4" gateway:

1 Open System Preferences and click Network.

2 Below the list of Interfaces, click the Add (+) button.

3 From the Interface pop-up menu, choose "6 to 4".

4 In the Service name field, enter a unique name for the service, then click Create.

5 If you have a relay address, choose Manually from the Configuration pop-up menu and enter it; otherwise, leave the Configuration pop-up menu set to Automatically.

6 Click Apply.

Where to Find More Information

The working group for the IPv6 website is www.ipv6.org.

A group of IPv6 enthusiasts maintains a list of applications that support IPv6 at www.ipv6forum.com.

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 technical details about a protocol in its RFC document.

You can search for RFC documents by number at www.ietf.org/rfc.html.

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

Below the "Accept recursive queries from the following networks" list, click the Add (+) button to add networks that recursive queries are accepted from, then enter the network address in the list.Below the "Forwarder IP Addresses" list, click the Add (+) button to add networks that unauthorized queries get forwarded to, then enter the network address in the list.

Glossary

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's 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 of either a **Certificate Authority** (CA) or the key user.

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

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 interacting with the computer (for example, to run programs or modify file system permissions) by entering text commands at a shell prompt. See also **shell**; **shell prompt**.

computer name The default name used for SLP and SMB service registrations. The Network Browser in the Finder uses SLP to find computers advertising Personal File Sharing and Windows File Sharing. It can be set to bridge subnets depending on the network router settings. When you turn on Personal File Sharing, users see the computer name in the Connect to Server dialog in the Finder. Initially it is "<first created user>'s Computer" (for example, "John's Computer") but can be changed to anything. The computer name is used for browsing for network file servers, print queues, Bluetooth® discovery, Apple Remote Desktop clients, and any other network resource that identifies computers by computer name rather than network address. The computer name is also the basis for the default local host name.

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 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 to control access to a server. A filter is made up of an IP address and a subnet mask, and sometimes a port number and access type. The IP address and the subnet mask determine the range of IP addresses that the filter applies to.

firewall Software that protects the network applications running on your server. IP firewall service, which is part of Mac OS X Server software, scans incoming IP packets and rejects or accepts these packets based on a set of filters you create.

forward zone The DNS zone that holds no records of its own, but forwards DNS queries to another zone.

FTP File Transfer Protocol. A protocol that allows computers to transfer files over a network. FTP clients using any operating system that supports FTP can connect to a file server and download files, depending on their access privileges. Most Internet browsers and a number of freeware applications can be used to access an FTP server.

gateway A network node that interfaces one network to another. Often, it refers to a computer that links a private LAN to a public WAN, with or without Network Address Translation (NAT). A router is a special kind of gateway that links related network segments.

GB Gigabyte. 1,073,741,824 (2^{30}) 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 computer, historically referred to as the UNIX hostname.

HTTP Hypertext Transfer Protocol. The client/server protocol for the World Wide Web. HTTP provides a way for a web browser to access a web server and request hypermedia documents created using HTML.

Hypertext 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 A set of interconnected computer networks communicating through a common protocol (TCP/IP). The Internet is the most extensive publicly accessible system of interconnected computer networks in the world.

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 provicer 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 data packets and TCP keeps track of data packets.

IP address A unique numeric address that identifies a computer on the Internet.

IP subnet A portion of an IP network, which may be a physically independent network segment, that shares a network address with other portions of the network and is identified by a subnet number.

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 e-commerce applications as well as mail services.

KB Kilobyte. 1,024 (2^{10}) 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 it resides on.

local hostname A name that designates a computer on a local subnet. It can be used without a global DNS system to resolve names to IP addresses. It consists of lowercase letters, numbers, or hyphens (except as the last characters), and ends with ".local" (For example, bills-computer.local). Although the default name is derived from the computer name, a user can specify this name in the Sharing pane of System Preferences. It can be changed easily, and can be used anywhere a DNS name or fully qualified domain name is used. It can only resolve on the same subnet as the computer using it.

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 method for VPN. This authentication method encodes passwords when they are sent over the network and stores them in a scrambled form on the server. It offers good security during network transmission. MS-CHAP is a proprietary version of CHAP.

multicast 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. Called Bonjour (previously Rendezvous) by Apple, this proposed Internet standard protocol is sometimes referred to as ZeroConf or multicast DNS. For more information, visit www.apple.com or www.zeroconf.org. To see how this protocol is used in Mac OS X Server, see **local hostname**.

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 An older Apple protocol 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**.

NFS Network File System. A client/server protocol that uses Internet Protocol (IP) to allow remote users to access files as though they were local. NFS can export shared volumes to computers based on IP address, and also supports single sign-on (SSO) authentication through Kerberos.

nfsd daemon An NFS server process that runs continuously behind the scenes and processes NFS protocol and mount protocol requests from clients. nfsd can have multiple threads. The more NFS server threads, the better concurrency.

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, Active Directory protocols, or 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 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.

permissions Settings that define the kind of access users have to shared items in a file system. You can assign four types of permissions to a share point, folder, or file: Read & Write, Read Only, Write Only, and No Access. See also **privileges**.

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.

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 is usually 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).

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 stratums 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 (for example, different floors of a building) or by usage (for example, all eighth-grade students). The use of subnets simplifies administration. See also **IP subnet**.

subnet mask A number used in IP networking to specify which portion of an IP address is the network number.

TCP Transmission Control Protocol. A method used with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. IP handles the actual delivery of the data, and TCP keeps track of the units of data (called packets) into which a message is divided for efficient routing through the Internet.

time server A network server whose clock other computers on the network synchronize their own clocks to, 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 on 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.

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 they rely on having the same encryption system at both ends. The encryption may be performed by firewall software or by routers.

WAN Wide area network. A network maintained across geographically separated facilities, as opposed to a LAN (local area network) within a facility. Your WAN interface is usually the one connected to the Internet.

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.

Index

E

F

G

H

I

J

K

L

M

N

T

U

V

W

X

Z