



Mac OS X Server Network Services Administration

For Version 10.3 or Later



 Apple Computer, Inc.
© 2003 Apple Computer, Inc. All rights reserved.

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

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

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

Adobe and PostScript are trademarks of Adobe Systems Incorporated.

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

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

034-2351/9-20-03

Contents

Preface	5	How to Use This Guide
	5	What's Included in This Guide
	5	Using This Guide
	6	Setting Up Mac OS X Server for the First Time
	6	Getting Help for Everyday Management Tasks
	6	Getting Additional Information
Chapter 1	7	DHCP Service
	7	Before You Set Up DHCP Service
	9	Setting Up DHCP Service for the First Time
	10	Managing DHCP Service
	14	Monitoring DHCP Service
	16	Where to Find More Information
Chapter 2	17	DNS Service
	18	Before You Set Up DNS Service
	18	Setting Up DNS Service for the First Time
	21	Managing DNS Service
	22	Managing Zones
	25	Managing Records
	28	Monitoring DNS
	30	Securing the DNS Server
	33	Common Network Administration Tasks That Use DNS Service
	37	Configuring BIND Using the Command Line
	41	Where to Find More Information
Chapter 3	43	IP Firewall Service
	45	Understanding Firewall Filters
	48	Setting Up Firewall Service for the First Time
	49	Managing Firewall Service
	55	Monitoring Firewall Service
	57	Practical Examples
	59	Common Network Administration Tasks That Use Firewall Service
	60	Advanced Configuration

	63	Port Reference
	66	Where to Find More Information
Chapter 4	67	NAT Service
	67	Starting and Stopping NAT Service
	68	Configuring NAT Service
	68	Monitoring NAT Service
	69	Where to Find More Information
Chapter 5	71	VPN Service
	72	VPN and Security
	73	Before You Set Up VPN Service
	73	Managing VPN Service
	76	Monitoring VPN Service
	77	Where to Find More Information
Chapter 6	79	NTP Service
	79	How NTP Works
	80	Using NTP on Your Network
	80	Setting Up NTP Service
	81	Configuring NTP on Clients
	81	Where to Find More Information
Chapter 7	83	IPv6 Support
	84	IPv6 Enabled Services
	84	IPv6 Addresses in the Server Admin
	84	IPv6 Addresses
	86	Where to Find More Information
Glossary	87	
Index	95	

How to Use This Guide

What's Included in This Guide

This guide consists primarily of chapters that tell you how to administer various Mac OS X Server network services:

- DHCP
- DNS
- IP Firewall
- NAT
- VPN
- NTP
- IPv6 Support

Using This Guide

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

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

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

Setting Up Mac OS X Server for the First Time

If you haven't installed and set up Mac OS X Server, do so now.

- Refer to *Mac OS X Server Getting Started for Version 10.3 or Later*, the document that came with your software, for instructions on server installation and setup. For many environments, this document provides all the information you need to get your server up, running, and available for initial use.
- Review this guide to determine which services you'd like to refine and expand, to identify new services you'd like to set up, and to learn about the server applications you'll use during these activities.
- Read specific chapters to learn how to continue setting up individual services. Pay particular attention to the information in these sections: "Setup Overview," "Before You Begin," and "Setting Up for the First Time."

Getting Help for Everyday Management Tasks

If you want to change settings, monitor services, view service logs, or do any other day-to-day administration task, you can find step-by-step procedures by using the on-screen help available with server administration programs. While all the network services' administration tasks are also documented in the network services administration guide, sometimes it's more convenient to retrieve information in onscreen help form while using your server.

Getting Additional Information

In addition to this document, you'll find information about Mac OS X Server:

- In *Mac OS X Server Getting Started for Version 10.3 or Later*, which tells you how to install and set up your server initially.
- At www.apple.com/server.
- In onscreen help on your server.
- In Read Me files on your server CD.

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

You can use the DHCP module in Server Admin to:

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

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

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

Before You Set Up DHCP Service

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

Creating Subnets

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

Assigning IP Addresses Dynamically

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

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 5, "VPN Service," on page 71.

Using Static IP Addresses

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

Static IP addresses must be set up manually by entering the IP address on the computer or device that is assigned the address. Manually configured static IP addresses avoid possible issues certain services may have with DHCP-assigned addresses and avoid the delay required for DHCP to assign an address.

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

Locating the DHCP Server

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

Interacting With Other DHCP Servers

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

Using Multiple DHCP Servers on a Network

You can have multiple DHCP servers on the same network. However, it's important that they're configured properly as to not interfere with each other. Each server needs a unique pool of IP addresses to distribute.

Assigning Reserved IP Addresses

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

Getting More Information on the DHCP Process

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

Setting Up DHCP Service for the First Time

If you used the Setup Assistant to configure ports on your server when you installed Mac OS X Server, some DHCP information is already configured. You need to follow the steps in this section to finish configuring DHCP service. You can find more information about settings for each step in "Managing DHCP Service" on page 10.

Step 1: Create subnets

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

See "Creating Subnets in DHCP Service" on page 10.

Step 2: Set up logs for DHCP service

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

DHCP service records diagnostic messages in the system log file. To keep this file from growing too large, you can suppress most messages by changing your log settings in the Logging pane of the DHCP service settings. For more information on setting up logs for DHCP service, see “Setting the Log Detail Level for DHCP Service” on page 15.

Step 3: Start DHCP service

See “Starting and Stopping DHCP Service” on page 10.

Managing DHCP Service

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

Starting and Stopping DHCP Service

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

To start or stop DHCP service:

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

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

Creating Subnets in DHCP Service

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

To create a new subnet:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Click Add, or double-click an existing subnet.
- 5 Select the General tab.
- 6 Enter a descriptive name for the new subnet. (Optional)

- 7 Enter a starting and ending IP address for this subnet range.
Addresses must be contiguous, and they can't overlap with other subnets' ranges.
- 8 Enter the subnet mask for the network address range.
- 9 Choose the Network Interface from the pop-up menu.
- 10 Enter the IP address of the router for this subnet.
If the server you're configuring now is the router for the subnet, enter this server's internal LAN IP address as the router's address.
- 11 Define a lease time in hours, days, weeks, or months.
- 12 If you wish to set DNS, LDAP, or WINS information for this subnet, enter these now.
See "Setting the DNS Server for a DHCP Subnet" on page 12, "Setting LDAP Options for a Subnet" on page 13, and "Setting WINS Options for a Subnet" on page 13 for more information.
- 13 Click Save.

Changing Subnet Settings in DHCP Service

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

To change subnet settings:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Subnets tab.
- 4 Select a subnet.
- 5 Click Edit.
- 6 Make the changes you want.
These changes can include adding DNS, LDAP, or WINS information. You can also redefine address ranges or redirect the network interface that responds to DHCP requests.
- 7 Click Save.

Deleting Subnets From DHCP Service

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

To delete subnets or address ranges:

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

Changing IP Address Lease Times for a Subnet

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

To change the lease time for a subnet address range:

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

Setting the DNS Server for a DHCP Subnet

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

To set DNS options for a subnet:

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

Setting LDAP Options for a Subnet

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

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

To set LDAP options for a subnet:

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

Setting WINS Options for a Subnet

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

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

To set WINS options for a subnet:

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

Disabling Subnets Temporarily

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

To disable a subnet:

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

Monitoring DHCP Service

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

Viewing the DHCP Status Overview

The status overview shows a simple summary of the DHCP service. It shows whether or not the service is running, how many clients it has, and when service was started. It also shows how many IP addresses are statically assigned from your subnets and the last time the client database was updated.

To see the overview:

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

Setting the Log Detail Level for DHCP Service

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

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

To set up the log detail level:

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

Viewing DHCP Log Entries

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

To see DHCP log entries:

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

Viewing the DHCP Client List

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

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

To view the DHCP client list:

- 1 In Server Admin, choose DHCP from the Computers & Services list.
- 2 Click Clients.
Click any column heading to sort the list by different criteria.

Where to Find More Information

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

For details about DHCP, see RFC 2131.

For more information on `bootpd` and its advanced configuration options, see `bootpd`'s man page.

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

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

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

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

- You don't have the option to use DNS from your ISP or other source.
- You plan on making frequent changes to the namespace and want to maintain it yourself.
- You have a mail server on your network and you have difficulties coordinating with the ISP that maintains your domain.

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

Before You Set Up DNS Service

This section contains information you should consider before setting up DNS on your network. The issues involved with DNS administration are complex and numerous. You should only set up DNS service on your network if you're an experienced DNS administrator.

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

DNS and BIND

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

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

Setting Up Multiple Name Servers

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

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

Setting Up DNS Service for the First Time

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

Step 1: Register your domain name

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

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

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

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

Step 2: Learn and plan

If you’re new to working with DNS, learn and understand DNS concepts, tools, and features of Mac OS X Server and BIND. See “Where to Find More Information” on page 41.

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

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

There are two ways to configure DNS service on Mac OS X Server. First, and recommended, you can use Server Admin to set up DNS service. For more information, see “Managing DNS Service” on page 21 for instructions.

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

The configuration file is located in this file:

`/etc/named.conf`

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

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

See “Configuring BIND Using the Command Line” on page 37 for more information.

Step 3: Configure basic DNS settings

See “Managing DNS Service” on page 21 for more information.

Step 4: Create a DNS Zone

Use Server Admin to set up DNS zones. See “Managing Zones” on page 22 for instructions. After adding a master zone, Server Admin automatically creates an NS record with the same name as the Source of Authority (SOA).

Step 5: Add Address and additional records to the zone.

Use Server Admin to add additional records to your Zone. Create an Address record for every computer or device (printer, file server, etc.) that has a static IP address and needs a name. When you create an A record, you have the option to specify the creation of a reverse lookup record and it’s corresponding zone. See “Managing Records” on page 25 for instructions.

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

If you provide mail service over the Internet, you need to set up an MX record for your server. See “Setting Up MX Records” on page 33 for more information.

Step 7: Configure the reverse lookup zone (optional)

For each zone that you create, Mac OS X Server creates a reverse lookup zone. Reverse lookup zones translate IP addresses to domain names, rather than normal lookups which translate domain names to IP addresses. If you have not specified reverse lookup records when initially creating your A records, you might need to configure your reverse lookup zone after its creation.

Step 8: Start DNS service

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

See “Starting and Stopping DNS Service” on page 21 for more information.

Managing DNS Service

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

Starting and Stopping DNS Service

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

To start or stop DNS service:

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

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

Enabling or Disabling Zone Transfers

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

To enable or disable zone transfer:

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

Enabling or Disabling Recursion

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

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

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

To enable or disable recursion:

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

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

Managing Zones

Zones are the basic organizational unit of the Domain Name System. Zones contain records and are defined by how they acquire those records, and how they respond to DNS requests. There are three kinds of zones:

Master

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

Slave

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

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

Forward

A forward zone directs all lookup requests for that zone to other DNS servers. Forward zones don't do zone transfers. Often, forward zone servers are used to provide DNS services to a private network behind a firewall. In this case, the DNS server must have access to the Internet and a DNS server outside the firewall.

Adding a Master Zone

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

To add a master zone:

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

The zone name must have a trailing period: "example.com."

- 6 Choose Master from the Zone Type pop-up menu.
- 7 Enter the hostname of the domain's SOA.

If this computer will be the authoritative name server for the domain, enter the computer's hostname (with a trailing period). For example, "ns.example.com."

- 8 Enter the email address of the zone's administrator.

The email address must not have an "@"; but a period; it should also have a trailing period. For example, the email address "admin@example.com" should be entered as "admin.example.com." (Remember to leave the trailing period.)

- 9 Click OK and then click Save.

Adding a Slave Zone

A slave zone is a copy of a master zone stored on a slave or secondary name server. Slaves must be configured to request the copy of the master zone's data. Slave zones use zone transfers to get copies of the master zone data.

To add a slave zone:

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

The Zone name must have a trailing period: "example.com."

- 6 Choose Slave from the Zone Type pop-up menu.
- 7 Click OK.
- 8 Click Add under the "Master servers for backup" pane.
- 9 Enter the IP addresses for the master servers for this zone.
- 10 Click Save.

Adding a Forward Zone

A forward zone directs all lookup requests to other DNS servers.

To add a forward zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Click Add beneath the Zones list.
- 5 Enter a zone name.
The Zone name must have a trailing period: "example.com."
- 6 Choose the Forward zone type from the Zone Type pop-up menu.
- 7 Click OK.
- 8 Click Add under the "Forward servers for fwd" pane.
- 9 Enter the IP addresses for the master servers for this zone.
- 10 Click Save.

Duplicating a Zone

You can create a copy of an existing zone on the same computer. You could use this to speed up configuration of multiple zones.

To duplicate a zone:

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

Modifying a Zone

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

To modify a zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Click the Edit button beneath the Zones list.
- 5 Change the zone name, type, or administrator email address as needed.
For more information on zone types, see "Managing Zones" on page 22.
- 6 Click OK, and click Save.

Deleting a Zone

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

To delete a zone:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Click the Delete button beneath the Zones list.
- 5 Click Save to confirm the deletion.

Managing Records

Each zone contains a number of records. These records are requested when a client computer needs to translate a domain name (like `www.example.com`) to an IP number. Web browsers, email clients, and other network applications rely on a zone's records to contact the appropriate server.

The master zone's records will be queried by others across the Internet so they can connect to your network services. There are several kinds of DNS records. The records which are available for configuration by Server Admin's user interface are:

- *Address (A)*: Stores the IP address associated with a domain name.
- *Canonical Name (CNAME)*: Stores the "real name" of a server when given a "nickname" or alias. For example, `mail.apple.com` might have a canonical name of `MailSrv473.apple.com`.
- *Mail Exchanger (MX)*: Stores the domain name of the computer that is used for email in a zone.

- *Name Server (NS)*: Stores the authoritative name server for a given zone.
- *Pointer (PTR)*: Stores the domain name of a given IP address (reverse lookup).
- *Text (TXT)*: Stores a text string as a response to a DNS query.

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

Adding a Record to a Zone

You need to add records for each domain name (example.com) and subdomain name (machine.example.com) for which the DNS master zone has responsibility. You should not add records for domain names that this zone doesn't control.

To add a record:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the Zone to which this record will be added.
- 5 Click the Add button beneath the Records list.
- 6 Select a record type from the Type pop-up menu.
- 7 In the first field, enter the fully qualified domain name.

The domain name must have a trailing period: "example.com."

If you're creating a PTR record, enter the IP address instead.

If you're creating a TXT record, enter the text string you want.

- 8 In the second field, for the following record types, enter:
 - *A records*: the IP address.
 - *AAAA records*: the IPv6 address.
 - *C-NAME records*: the real name of the computer.
 - *MX records*: the name (with trailing period) or IP address of the domain's mail exchanger.
 - *PTR records*: the full domain name with trailing period.
- 9 If creating an A record, select "Create reverse mapping record" to automatically create its corresponding PTR record.
- 10 Click OK, and click Save.

Modifying a Record in a Zone

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

To modify a record:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the Zone in which this record will be modified.
- 5 Double-click the record to be modified, or select the record and click the Edit button.
- 6 Modify the record as needed.
You can change the hostname, record type, or IP number.
- 7 Click OK.

Deleting a Record From a Zone

You should delete records whenever a domain name is no longer associated with a working address.

To delete a record:

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

Monitoring DNS

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

Viewing DNS Service Status

You can check the DNS Status window to see:

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

To view DNS service status:

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

Viewing DNS Service Activity

You can check the DNS Status window to see:

- The number of transfers running and deferred.
- Whether the service is loading the configuration file.
- If the service is priming.
- Whether query logging is turned on or off.
- The number of Start of Authority (SOA) queries in progress.

To view DNS service activity:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Activity to view operations currently in progress.

Viewing DNS Log Entries

DNS service creates entries in the system log for error and alert messages.

To see DNS log entries:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Log.

Changing DNS Log Detail Levels

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

To change the log detail level:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.
- 4 Choose the detail level from the Log Level pop-up menu.

The possible log levels are:

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

Changing DNS Log File Location

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

To change the log detail level:

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

If no path is entered, the default location is `/var/logs/`.

Viewing DNS Usage Statistics

You can check the DNS Statistics window to see statistics on common DNS queries.

Some common DNS queries begin with the following:

- *Name Server (NS)*: Asks for the authoritative name server for a given zone.
- *Address (A)*: Asks for the IP address associated with a domain name.
- *Canonical Name (CName)*: Asks for the “real name” of a server when given a “nickname” or alias. For example, `mail.apple.com` might have a canonical name of `MailSrv473.apple.com`.
- *Pointer (PTR)*: Asks for the domain name of a given IP address (reverse lookup).
- *Mail Exchanger (MX)*: Asks which computer in a zone is used for email.
- *Start Of Authority (SOA)*: Asks for name server information shared with other name servers and possibly the email address of the technical contact for this name server.
- *Text (TXT)*: Asks for text records used by the administrator.

To see DNS usage statistics:

- 1 In Server Admin, choose DNS in the Computer & Services list.
- 2 Click Activity to view operations currently in progress and usage statistics.

Securing the DNS Server

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

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

DNS Spoofing

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

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

Also, a falsified mail record could allow a hacker to intercept mail sent to or from a domain. If the hacker also forwards those emails to the correct mail server after copying them, this can go undetected indefinitely.
- Prevent proper domain name resolution and access to the Internet.
This is the most benign of DNS spoof attacks. It merely makes a DNS server appear to be malfunctioning.

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

Server Mining

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

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

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

To specify zone transfer IP addresses:

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

Follow the instructions in “Creating an Advanced IP Filter for TCP ports” in Chapter 3, “IP Firewall Service.” Use the following settings:

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

DNS Service Profiling

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

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

To alter BIND’s version response:

- 1 Launch a command-line text editor (like vi, emacs, or pico).
- 2 Open `named.conf` for editing.
- 3 Add the following to the “options” brackets of the configuration file.
`version "[your text, maybe 'we're not telling!']";`
- 4 Save the config file.

Denial-of-Service (DoS)

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

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

The easiest way to guard against this attack is to block the offending IP address with your firewall. See “Creating an Advanced IP Filter for TCP ports” on page 51. Unfortunately, this means the attack is already underway and the hacker’s queries are being answered and the activity logged.

Service Piggybacking

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

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

To prevent recursion entirely, see “Enabling or Disabling Recursion” on page 21.

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

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

Please see BIND’s documentation for further information.

Common Network Administration Tasks That Use DNS Service

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

Setting Up MX Records

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

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

user-name@reliable.example.com

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

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

example.com

10 reliable.example.com

20 our-backup.example.com

30 last-resort.example.com

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

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

Configuring DNS for Mail Service

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

To enable MX records:

- 1 In Server Admin, choose DNS in the Computers & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the Zone you want to use.
- 5 Click the Add button under the Records pane.
- 6 Choose MX from the Type pop-up menu.
- 7 Enter the domain name (like "example.com.") in the From field.
- 8 Enter the name of the mail server (like "mail.example.com.") in the To field.
- 9 Enter a precedence number.
- 10 Click OK.

Enabling Redundant Mail Servers

You may need to set up multiple servers for redundancy. If this is the case, you'll need to add additional information to each MX record. Create one record for each auxiliary server. This consists of two steps:

These instructions assume you have an existing MX record for a primary mail server. If not, please see "Configuring DNS for Mail Service" on page 34.

Step 1: Edit the MX record of the primary mail server

- 1 In Server Admin, choose DNS in the Computers & Services list.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the Zone you want to use.
- 5 Click the primary mail server's MX record in the Records pane.
- 6 Click the Edit button below the Records pane.
- 7 Enter a low precedence number for that server.

A lower number indicates it will be chosen first, if available, to receive mail.

- 8 Click OK.
- 9 Proceed to Step 2.

Step 2: Create records and priorities for the auxiliary mail servers

These instructions assume you have edited the original MX record. If not, please do so before proceeding.

These instructions also assume you have already set up and configured one or more auxiliary mail servers.

To enable backup or redundant mail servers:

- 1 In Server Admin, select DNS in the Computers & Services pane.
- 2 Click Settings.
- 3 Select the Zones tab.
- 4 Select the Zone you want to use.
- 5 Click the Add button under the Records pane.
- 6 Choose MX from the Type pop-up menu.
- 7 Enter the domain name (like 'example.com.') in the From field.
- 8 Enter the name of the mail server (like 'backup.example.com.') in the To field.
- 9 Enter a precedence number for that server which is higher than that of the primary server.
A higher the number indicates it will be chosen if the primary server is unavailable.
- 10 Click OK.

Setting Up Namespace Behind a NAT Router

If you're behind a Network Address Translation (NAT) router, you have a special set of IP addresses that are only usable within the NAT environment. If you were to assign a domain name to these addresses outside of the NAT router, none of your domain names would resolve to the correct computer. See Chapter 4, "NAT Service," on page 67 for more information about NAT.

You can, however, run a DNS service behind the router, assigning host names to the NAT IP addresses. This way, if you're behind the NAT router, 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 router to allow resolution of names outside the routed area. Your clients' networking settings should specify the DNS server behind the NAT router. The process of setting up one of these networks is the same as setting up a private network. See "Setting Up a Private TCP/IP Network" on page 36 for more information.

If you choose to do this, names entered by users outside the NAT router won't resolve to the addresses behind it. You should set the DNS records outside the NAT-routed area to point to the NAT router, and use NAT port forwarding to access computers behind the NAT router. For more information on port forwarding, see Chapter 4, "NAT Service," on page 67.

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

Network Load Distribution (aka Round Robin)

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

You enable round robin by adding multiple address entries in your zone data file for a given host. For example, suppose you want to distribute web server traffic between three servers on your network that all mirror the same content. Suppose the servers have the IP addresses 192.168.12.12, 192.168.12.13, and 192.168.12.14. You would add these lines to the zone data file db.example.com:

```
www.example.com 60 IN A 192.168.12.12
www.example.com 60 IN A 192.168.12.13
www.example.com 60 IN A 192.168.12.14
```

When BIND encounters multiple entries for one host, its default behavior is to answer queries by sending out this list in a cycled order. The first request gets the addresses in the order A, B, C. The next request gets the order B, C, A, then C, A, B, and so on. Notice that the *time-to-live* (TTL) in the second column is set quite short to mitigate the effects of local caching.

Setting Up a Private TCP/IP Network

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

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

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

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

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

Hosting Several Internet Services With a Single IP Address

You must have one server supplying all your Internet services (like mail, web). They may all be running on one computer with a single IP address. For example, you may want to have the domain name `www.example.com` resolve to the same IP address as `ftp.example.com`, or `mail.example.com`.

Setting up the DNS records for this service is easy. You'll still need a full set of DNS records, one for each name you want to resolve.

- Set up MX records for mail, so `mail.example.com` resolves to your server's IP address.
- Set up A records for each service your server provides, so `web.example.com` resolves to your server's IP address.
- Do the same for each service you provide (`ftp.apple.com`, or `fileshare.apple.com`, or whatever).

As your needs grow, you can add other computers to the network to take over these services. Then all you have to do is update the DNS record, and your client's settings can remain the same.

Configuring BIND Using the Command Line

In order to set up and use DNS service on Mac OS X Server you may wish to configure BIND with the command-line. Configuring BIND requires making changes to UNIX configuration files in the Terminal application. To configure BIND, you must be comfortable with typing UNIX commands and using a UNIX text editor. Only manipulate these settings if you have a thorough understanding of DNS and BIND, preferably as an experienced DNS administrator.

Warning: Incorrect BIND configurations can result in serious network problems.

What Is BIND?

BIND stands for Berkeley Internet Name Domain. BIND runs on UNIX-based operating systems and is distributed as open-source software. BIND is used on the majority of name servers on the Internet today.

BIND is configured by editing text files containing information about how you want BIND to behave and information about the servers on your network. If you wish to learn more about DNS and BIND, resources are listed at the end of this chapter.

BIND on Mac OS X Server

Mac OS X Server uses BIND version 9.2.2. You can start and stop DNS service on Mac OS X Server using the Server Admin application. You can use Server Admin to view DNS status and usage statistics.

BIND Configuration File

By default, BIND looks for a configuration file labeled “named.conf” in the /etc directory. This file contains commands you can use to configure BIND’s many options. It also specifies the directory to use for zone data files.

Zone Data Files

Zone data files consist of paired address files and reverse lookup files. Address records link host names (host1.example.com) to IP addresses. Reverse lookup records do the opposite, linking IP addresses to host names. Address record files are named after your domain name— for example, example.com. Reverse lookup file names look like part of an IP address, such as db.192.168.12.

By default, the zone data files are located in /var/named/.

Practical Example

The following example allows you to create a basic DNS configuration using BIND for a typical network behind a Network Address Translation (NAT) device that connects to an ISP. The port (cable modem/DSL/dial-up/etc.) that is connected to your ISP is referred to here as the *WAN interface*. The port that is connected to your internal network is referred to here as the *LAN interface*. The sample files you need are installed with Mac OS X Server in the directories listed in the steps below. This example also assumes the following:

- The IP address of the WAN interface is determined by your ISP.
- The IP address of the LAN interface is 10.0.1.1.
- The IP address of the Mac OS X or Mac OS X Server computer that will be used as the DNS server is 10.0.1.2.
- The IP addresses for client computers are 10.0.1.3 through 10.0.1.254.

If IP address assignment is provided by the NAT device via DHCP, it must be configured with the above information. Please consult your router or gateway manual for instructions on configuring its DHCP server.

If your NAT device connects to the Internet, you also need to know the DNS server addresses provided by your ISP.

Setting Up Sample Configuration Files

The sample files can be found in `/usr/share/named/examples`.

The sample files assume a domain name of `example.com` behind the NAT. This may be changed, but must be changed in *all* modified configuration files. This includes renaming `/var/named/example.com.zone` to the given domain name, for example, `/var/named/foo.org.zone`

To set up the sample files:

- 1 In Terminal, log in as root.

- 2 Enter the following command:

```
cp /etc/named.conf /etc/named.conf.OLD
```

This saves a copy of the process configuration file.

- 3 Then, enter the following command:

```
cp /usr/share/named/examples/db.10.0.1.sample /var/named/10.0.1.zone
```

This copies the sample file for the NAT zone.

- 4 Enter the following command:

```
cp /usr/share/named/examples/example.com.sample /var/named/  
example.com.zone
```

This copies the sample file for your domain.

- 5 Now, enter the following command:

```
cp /usr/share/named/examples/named.conf.sample /etc/named.conf
```

This copies in a sample named process configuration file.

- 6 Using a command-line text editor (like `pico`, or `emacs`), open `/etc/named.conf` for editing.

- 7 Follow the instructions in the sample file to apply edits appropriate to your specific installation.

- 8 Save your changes to `named.conf`.

- 9 Use Server Admin to start DNS service.

- 10 In the Network pane of System Preferences, change the domain name servers to list only the IP address of the new DNS server, `10.0.1.2`.

Configuring Clients

If the IP addresses of your client computers are statically assigned, change the domain name servers of their Network preference panes to only list the new server's IP address, `10.0.1.2`.

If you are using Mac OS X Server as your DHCP Server:

- 1 In Server Settings, click the Network tab, click DHCP/NetBoot, and choose Configure DHCP/NetBoot.
- 2 On the Subnet tab, select the subnet on the built-in Ethernet port and click Edit.
- 3 In the General tab, enter the following information:

Start: 10.0.1.3

End: 10.0.1.254

Subnet Mask: 255.255.255.0

Router: 10.0.1.1

- 4 Click the DNS tab and enter the following information:

Default Domain: example.com

DNS Servers: 10.0.1.2

- 5 Click the Save button and log out of Server Settings.

Note: The client computers may not immediately populate with the new IP configuration information. This will depend upon when their DHCP leases expire. It may be necessary to restart the client computers for the changes to populate.

Check Your Configuration

To verify the steps were successful, open Terminal, located in /Applications/Utilities and enter the following commands (substituting the local domain name for "server.example.com" as appropriate):

```
dig server.example.com
```

```
dig -x 10.0.1.2
```

Note: If this generic configuration example does not meet your needs, Apple recommends that you don't attempt to configure DNS on your own and that you seek out a professional consultant or additional documentation.

Using DNS With Dynamically Assigned IP Addresses

Dynamic DNS is a mechanism that lets you modify the IP address/domain name list without directing the name server to reload the edited list. This means you can update the name server remotely and easily modify DNS data.

You can use dynamic DNS with DHCP service. DHCP assigns each client computer a dynamic IP address when the computer starts up. Because a DHCP server may assign IP addresses randomly, it can be useful to assign meaningful DNS names to these addresses on the fly.

For instance, if “Bob” walks into work in the morning and starts up his computer, and the DHCP server assigns his computer a dynamic IP address, a DNS entry “bob.example.com” can be associated with that IP address. Even though Bob’s IP address may change every time he starts up his computer, his DNS name remains the same. This lets users communicate with Bob’s computer without knowing the IP address.

You can also use dynamic DNS to provide static host names for users who connect to the Internet through a modem. An ISP can set up dynamic DNS so a home computer has the same host name every time it connects.

Where to Find More Information

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

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


Request For Comment Documents

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

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

Firewall service is software that protects the network applications running on your Mac OS X Server. Turning on firewall service is similar to erecting a wall to limit access. Firewall service scans incoming IP packets and rejects or accepts these packets based on the set of filters you create. You can restrict access to any IP service running on the server, and you can customize filters for all incoming clients or for a range of client IP addresses.

The illustration below shows an example firewall process.



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

- If the port number is in the filter list, the filter applied is the one that contains the most specific address range.
- If the port number is not in the list, the Default filter that contains the most specific address range is used.

The port filters you create are applied to TCP packets and can also be applied to UDP packets. In addition, you can set up filters for restricting Internet Control Message Protocol (ICMP), Internet Group Management Protocol (IGMP), and NetInfo data.

Important: When you start firewall service the first time, most all incoming TCP packets are denied until you change the filters to allow access. By default, only the ports essential to remote administration are available. These include access by Remote Directory Access (625), Server Administration via Server Admin (687), and Secure Shell (22). For any other network service, you must create filters to allow access to your server. If you turn firewall service off, all addresses are allowed access to your server.

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

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

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

Mac OS X Server uses the application ipfw for firewall service.

Understanding Firewall Filters

When you start firewall service, the default configuration denies access to all incoming packets from remote computers except ports for remote configuration. This provides a high level of security. You can then add new IP filters to allow server access to those clients who require access to services.

To learn how IP filters work, read the following section. To learn how to create IP filters, see “Managing Firewall Service” on page 49.

What is a Filter?

A filter is made up of an IP address and a subnet mask, and sometimes a port number and access type. The IP address and the subnet mask together determine the range of IP addresses to which the filter applies, and can be set to apply to all addresses.

IP Address

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

Subnet Mask

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

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

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

CIDR	Corresponds to Netmask	Number of addresses in the range
/1	128.0.0.0	4.29x10 ⁹
/2	192.0.0.0	2.14x10 ⁹
/3	224.0.0.0	1.07x10 ⁹
/4	240.0.0.0	5.36x10 ⁸
/5	248.0.0.0	1.34x10 ⁸
/6	252.0.0.0	6.71x10 ⁷
/7	254.0.0.0	3.35x10 ⁷
/8	255.0.0.0	1.67x10 ⁷
/9	255.128.0.0	8.38x10 ⁶
/10	255.192.0.0	4.19x10 ⁶
/11	255.224.0.0	2.09x10 ⁶
/12	255.240.0.0	1.04x10 ⁶
/13	255.248.0.0	5.24x10 ⁵
/14	255.252.0.0	2.62x10 ⁵
/15	255.255.0.0	1.31x10 ⁵
/16	255.255.255.0	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 filters using Server Admin, you enter an IP address and the CIDR format subnet mask. Server Admin shows you the resulting address range, and you can change the range by modifying the subnet mask. When you indicate a range of possible values for any segment of an address, that segment is called a *wildcard*. The following table gives examples of address ranges created to achieve specific goals.

Goal	Sample IP address	Enter this in the address field:	Address range affected
Create a filter 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 filter 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 filter 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 filter that applies to all incoming addresses.		Select "Any"	All IP addresses

Rule Mechanism and Precedence

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

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

Multiple IP Addresses

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

Setting Up Firewall Service for the First Time

Once you've decided which filters you need to create, follow these overview steps to set up firewall service. If you need more help to perform any of these steps, see "Managing Firewall Service" on page 49 and the other topics referred to in the steps.

Step 1: Learn and plan

If you're new to working with IP Firewall, learn and understand firewall concepts, tools, and features of Mac OS X Server and BIND. For more information, see "Understanding Firewall Filters" on page 45.

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

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

Step 2: Start firewall service

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

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

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

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

You can create additional groups based on source IP number or destination IP number. See "Creating an Address Group" on page 50 for more information.

Step 4: Add filters to the IP filter list

Read "Understanding Firewall Filters" on page 45 to learn how IP filters work and how to create them. You use this to further all other services, strengthen your network security, and manage your network traffic through the firewall.

For information about creating a new filter, see "Creating an Advanced IP Filter for TCP ports" on page 51.

Step 5: Save firewall service changes

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

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

Managing Firewall Service

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

Starting and Stopping Firewall Service

By default, firewall service blocks all incoming TCP connections and allows all UDP connections. Before you turn on firewall service, make sure you've set up filters allowing access from IP addresses you choose. Otherwise, no one will have access to your server.

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

To start or stop firewall service:

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

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

Opening the Firewall for Standard Services

By default, firewall service blocks all incoming TCP connections and allows all UDP connections. Before you turn on firewall service, make sure you've set up filters allowing access from IP addresses you choose; otherwise, no one will have access to your server.

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

- Web service
- Apple File service
- Windows File service
- FTP service
- Printer Sharing

- DNS/Rendezvous
- ICMP Echo Reply (incoming pings)
- IGMP (Internet Gateway Multicast Protocol)
- PPTP VPN
- L2TP VPN
- QTSS media streaming
- iTunes Music Sharing

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

To open the firewall for standard services:

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

If you want to restrict or designate IP addresses for a standard service, you should create an address group rather than use the Any address group. See “Creating an Address Group” on page 50 for more information.

- 5 Select the services you want to allow.
- 6 Click Save.

Creating an Address Group

You can define groups of IP addresses for your firewall filters. These groups are used to organize and target the filters. The default address group is for all addresses.

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

To create an address group:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Click Add beneath the Address Group pane.
- 5 Enter a group name.
- 6 Enter the addresses and subnet mask you want the filters to effect.
- 7 Click OK.

Editing or Deleting an Address Group

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

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

To edit or delete an address group:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the group name from the Address Group pane.
- 5 Click the Edit button beneath the Address Group pane to edit it.
Click the Delete button beneath the Address Group pane to delete it.
- 6 Edit the Group name or addresses as needed.
- 7 Click OK.
- 8 Click Save.

Duplicating an Address Group

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

To duplicate an address group:

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

Creating an Advanced IP Filter for TCP ports

You can use the Advanced Settings pane to configure very specific filters for TCP ports. IP filters contain an IP address and a subnet mask. You can apply a filter to all IP addresses, a specific IP address, or a range of IP addresses.

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

To create an IP filter for TCP ports:

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

Alternatively, you can select a rule similar to the one you want to create, and click Duplicate then Edit.
- 5 Select whether this filter will allow or deny access in the Action pop-up menu.
- 6 Choose TCP from the Protocol pop-up menu.
- 7 Choose a TCP service from the pop-up menu.

If you want to select a nonstandard service port, choose Other.
- 8 If desired, choose to log packets that match the filter.
- 9 Enter the Source IP address range you want to filter.

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

If you have selected a nonstandard service port, enter the source port number.
- 10 Enter the Destination IP address range you want to filter.

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

If you have selected a nonstandard service port, enter the source port number.
- 11 Choose which network interface this filter applies to.
- 12 Click OK.
- 13 Click Save to apply the filter immediately.

Creating an Advanced IP Filter for UDP Ports

You can use the Advanced Settings pane to configure very specific filters for UDP ports. Many services use User Datagram Protocol (UDP) to communicate with the server. By default, all UDP connections are allowed. You should apply filters to UDP ports sparingly, if at all, because “deny” filters could create severe congestion in your server traffic.

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

You should also allow UDP port access for specific services, including:

- DNS
- DHCP
- SLP
- Windows Name Service browsing

- Remote Desktop
- NFS
- NetInfo

UDP ports above 1023 are allocated dynamically by certain services, so their exact port numbers may not be determined in advance.

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

To easily configure UDP access for these ports, see “Opening the Firewall for Standard Services” on page 49. If you need more advanced firewall settings for these basic UDP services, use the following instructions to create them.

To create an IP filter for UDP ports:

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

Alternatively, you can select a rule similar to the one you want to create, and click Duplicate then Edit.

- 5 Select whether this filter will allow or deny access in the Action pop-up menu.
- 6 Choose UDP from the Protocol pop-up menu.
- 7 Choose a UDP service from the pop-up menu.

If you want to select a nonstandard service port, choose Other.

- 8 If desired, choose to log packets that match the filter.
- 9 Enter the Source IP address range you want to filter.

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

If you have selected a nonstandard service port, enter the source port number.

- 10 Enter the Destination IP address range you want to filter.

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

If you have selected a nonstandard service port, enter the source port number.

- 11 Choose which network interface this filter applies to.
- 12 Click OK.
- 13 Click Save to apply the filter immediately.

Editing Advanced IP Filters

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

To edit advanced IP filters:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Advanced tab.
- 4 Select a filter and click Duplicate, Edit, or Delete. If you're deleting a filter, you've finished.
- 5 Make any changes to the settings, then click Save.

Changing the Default Filter

If the server receives a packet using a port or IP address to which none of your filters apply, firewall service uses the Default filter. You can set the Default filter to either deny or allow these packets for specific IP addresses. By default the Default filter denies access.

If you need to change the Default filter to allow access, you can. However, you shouldn't take this action lightly. Changing the default to allow means you must explicitly deny access to your services by setting up specific port filters for all the services that need protection.

It is recommended that you leave the Default filter in place and use the General panel to create higher priority rules which allow access to designated services.

To change the Default setting:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select Default and click Edit.
- 5 Make any changes to the settings, then Click Save.

Monitoring Firewall Service

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

Viewing the Firewall Status Overview

The Status Overview shows a simple summary of the firewall service. It shows whether or not the service is running and which filter rules are active.

To see the overview:

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

Setting Up Logs for Firewall Service

You can log only the packets that are denied by the filters you set, only the packets that are allowed, or both. Both logging options can generate a lot of log entries, which can fill up disk space and degrade the performance of the server. You should use "Log all allowed packets" only for limited periods of time.

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

To set up logs:

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

Viewing the Firewall Log

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

To view the log for firewall service:

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

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

Log Example 1

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

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

Log Example 2

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

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

Log Example 3

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

This entry shows that firewall service used rule 10 to send a packet to itself on port 660 via the loopback device 0.

Viewing Denied Packets

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

To view denied packets:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.
- 4 Make sure “Log denied packets” is checked.
- 5 View log entries in Server Admin by clicking Log.

Viewing Packets Logged by Filter Rules

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

To view filtered packets:

- 1 Turn on logging of filtered packets in filter editing window.
See “Editing Advanced IP Filters” on page 54 if you have not turned on logging for a particular filter.
- 2 To view log entries in Server Admin, choose Firewall from the Computers & Services list.
- 3 Click Log.

Practical Examples

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

Block Access to Internet Users

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

To do this:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the Any address group.
- 5 Make sure that Web Service is disabled in the right pane.
- 6 Click the Add button to create an address range.
- 7 Name the address group.
- 8 Add the local network address range.

This is done by using an example address from the network with its network mask in CIDR notation. For example, if a user has an address of 192.168.1.20 and the network mask is 255.255.255.0, then enter 192.168.1.20/24.

- 9 Click OK.
- 10 Select your newly created address group.
- 11 Select "Web Service" in the right pane to enable web access.
- 12 Click Save.

Block Junk Mail

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

To do this:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the Any address group.
- 5 Enable “SMTP Mail” in the right pane.
- 6 Click the Add button to create an address range.
- 7 Name the address group.
- 8 Enter 17.128.100.0 to the address range to indicate the junk mail sender’s address.
- 9 Click OK.
- 10 Select your newly created address group.
- 11 Deselect “SMTP Mail” in the right pane to disable mail transfer.
- 12 Click Save.

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

Allow a Customer to Access the Apple File Server

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

To do this:

- 1 In Server Admin, choose Firewall from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select the Any address group.
- 5 Disable “Apple File Service” in the right pane.
- 6 Click the Add button to create an address range.
- 7 Name the address group.
- 8 Enter 10.221.41.33 to the address range to indicate the customer’s address.
- 9 Click OK.
- 10 Select your newly created address group.
- 11 Select “Apple File Service” in the right pane to enable file access.
- 12 Click Save.

Common Network Administration Tasks That Use Firewall Service

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

Preventing Denial-of-Service (DoS) Attacks

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

To prevent ping denial-of-service attacks:

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

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

Controlling or Enabling Peer-to-Peer Network Usage

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

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

You can choose to limit P2P network usage to IP addresses behind the firewall. To do so, you'll need to open the P2P port for your LAN interface, but continue to block the port on the interface connected to the Internet (WAN interface). To learn how to make a firewall filter, see “Creating an Advanced IP Filter for TCP ports” on page 51.

Controlling or Enabling Network Game Usage

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

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

You can choose to limit network game usage to IP addresses behind the firewall. To do so, you'll need to open the appropriate port on your LAN interface, but continue to block the port on the interface connected to the Internet (WAN interface). Some games require a connection to a gaming service for play, so this may not be effective. To learn how to make a firewall filter, see "Creating an Advanced IP Filter for TCP ports" on page 51.

You can open the firewall to certain games, allowing network games to connect to other players and game services outside the firewall. To do this, you'll need to open up the appropriate port on your LAN and WAN interface. Some games require more than one port to be open. Consult the game's documentation for networking details. To learn how to make a firewall filter, see "Creating an Advanced IP Filter for TCP ports" on page 51.

Advanced Configuration

You might prefer to use a command-line interface and conventional configuration file to configure Mac OS X Server's firewall service. For example, you might have an existing ipfw configuration file that you want to migrate to a new Mac OS X Server installation. Alternately, you might need greater control of the firewall for troubleshooting or intrusion detection.

Background

When you click the Save button in Server Admin, all the old rules are flushed and new rules are loaded and apply immediately. This happens whether the IP firewall service is started or stopped. If the IP firewall service is running, it is stopped long enough to reload the rules, and it automatically restarts. The new rules are loaded from three sources:

- The rules from both the General and the Advanced panels (stored in `/etc/ipfilter/ip_address_groups.plist`).
- The manually configured ipfw rules, if any (stored in `/etc/ipfilter/ipfw.conf`).
- The NAT divert rule, if the NAT service is running.

If you want to put your own rules in the `ipfw.conf` file, you can use a template that is installed at `/etc/ipfilter/ipfw.conf.default`. Duplicate the file, rename it, and edit it as indicated in the template's comments.

Precautions

By using the Advanced panel or creating your own rules, you can put the server in a state that is completely cut off from network access. This might require a reboot in single-user-mode to restore network access. To avoid this, consider adding a cron job to disable the firewall periodically while you are testing rules. Be sure to disable this cron job when the machine is put into production.

The following command disables the firewall:

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

And this enables it:

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

Neither of these operations change the rules loaded into the firewall, they just determine whether those rules are applied.

Creating IP Filter Rules Using ipfw

You can use the `ipfw` command in conjunction with the firewall module of Server Admin when you want to:

- Display rules created by the firewall module. Each filter translates into one or more rules.
- Create filters with characteristics that can't be defined using the firewall module. For example, you may want to use rules specific to a particular kind of IP protocol. Or you may want to filter or block outgoing packets.
- Count the number of times rules are applied.

If you use `ipfw`, make sure you don't modify rules created using the firewall module. Changes you make to firewall module rules are not permanent. Firewall service recreates any rules defined using the firewall module whenever the service is restarted. Here is a summary of how the firewall module assigns rule numbers:

Rule number	Used by firewall module for
10	Loop back.
20	Discarding any packet from or to 127.0.0.0/8 (broadcast).
30	Discarding any packet from 224.0.0.0/3 (broadcast).
40	Discarding TCP packets to 224.0.0.0/3 (broadcast).
100–64000	User-defined port-specific filters.
63200	Denying access for icmp echo reply. Created when "Deny ICMP echo reply" is selected in the Advanced pane of the Configure Firewall window.

Rule number	Used by firewall module for
63300	Denying access for igmp. Created when Deny IGMP is selected in the Advanced pane of the Configure Firewall window.
63400	Allowing any TCP or UDP packet to access port 111 (needed by NetInfo). Created when a shared NetInfo domain is found on the server.
63500	Allowing user-specified TCP and UDP packets to access ports needed for NetInfo shared domains. You can configure NetInfo to use a static port or to dynamically select a port from 600 through 1023. Then use the Configure Firewall window to allow all or specific clients to access those ports.
64000–65000	User-defined filters for Default.

Reviewing IP Filter Rules

To review the rules currently defined for your server, use the Terminal application to submit the `ipfw show` command. The show command displays four columns of information:

Column	Information
1	The rule number. The lower the number, the higher the priority of the rule.
2	The number of times the filter has been applied since it was defined.
3	The number of bytes to which the filter has been applied.
4	A description of the rule.

When you type:

```
ipfw show
```

You will see information similar to this:

```
0010 260      32688    allow log ip from any to any via lo*
0020 0         0         deny log ip from 127.0.0.0/8 to any in
0020 0         0         deny log ip from any to 127.0.0.0/8 in
0030 0         0         deny log ip from 224.0.0.0/3 to any in
0040 0         0         deny log tcp from any to 224.0.0.0/3 in
001001       52         allow log tcp from 111.222.33.3 to 111.222.31.3 660
      in
...
```

Creating IP Filter Rules

To create new rules, use the `ipfw add` command. The following example defines rule 200, a filter that prevents TCP packets from a client with IP address 10.123.123.123 from accessing port 80 of the system with IP address 17.123.123.123:

```
ipfw add 200 deny tcp from 10.123.123.123 to 17.123.123.123 80
```

Deleting IP Filter Rules

To delete a rule, use the `ipfw delete` command. This example deletes rule 200:

```
ipfw delete 200
```

For more information, consult the man pages for `ipfw`.

Port Reference

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

TCP port	Used for	Reference
7	echo	RFC 792
20	FTP data	RFC 959
21	FTP control	RFC 959
22	ssh (secure shell)	
23	Telnet	RFC 854
25	SMTP (email)	RFC 821
53	DNS	RFC 1034
79	Finger	RFC 1288
80	HTTP (Web)	RFC 2068
88	Kerberos	RFC 1510
106	Open Directory Password Server (along with 3659)	
110	POP3 (email)	RFC 1081
111	Remote Procedure Call (RPC)	RFC 1057
113	AUTH	RFC 931
115	sftp	
119	NNTP (news)	RFC 977
123	Network Time Server synchronization (NTP)	RFC 1305
137	Windows Names	
138	Windows Browser	
139	Windows file and print (SMB)	RFC 100
143	IMAP (email access)	RFC 2060

TCP port	Used for	Reference
311	AppleShare IP remote Web administration, Server Monitor, Server Admin (servermgrd), Workgroup Manager (DirectoryService)	
389	LDAP (directory) Sherlock 2 LDAP search	RFC 2251
427	SLP (service location)	
443	SSL (HTTPS)	
514	shell	
515	LPR (printing)	RFC 1179
532	netnews	
548	AFP (AppleShare)	
554	Real-Time Streaming Protocol (QTSS)	RFC 2326
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)	
636	LDAP SSL	
660	Server Settings, Server Manager	
687	AppleShare IP Shared Users and Groups, Server Monitor, Server Admin (servermgrd)	
749	Kerberos administration using the kadmind command-line tool	
1220	QTSS Admin	
1694	IP Failover	
1723	PPTP VPN	RFC 2637
2049	NFS	
2236	Macintosh Manager	
3031	Program Linking	
3659	Open Directory Password Server (along with 106)	
7070	Real-Time Streaming Protocol (QTSS)	

TCP port	Used for	Reference
8000–8999	Web service	
16080	Web service with performance cache	
UDP port	Used for	Reference
7	echo	
53	DNS	
67	DHCP server (BootP)	
68	DHCP client	
69	Trivial File Transfer Protocol (TFTP)	
111	Remote Procedure Call (RPC)	
123	Network Time Protocol	RFC 1305
137	Windows Name Service (WINS)	
138	Windows Datagram Service	
161	Simple Network Management Protocol (SNMP)	
427	SLP (service location)	
497	Retrospect	
513	who	
514	Syslog	
554	Real-Time Streaming Protocol (QTSS)	
600–1023	Mac OS X RPC-based services (for example, NetInfo)	
985	NetInfo (when a shared domain is created using NetInfo Domain Setup)	
2049	Network File System (NFS)	
3031	Program Linking	
3283	Apple Network Assistant, Apple Remote Desktop	
5353	Rendezvous (mDNSResponder)	
6970 and up	QTSS	
7070	Real-Time Streaming Protocol alternate (QTSS)	

Where to Find More Information

For more information about ipfw:

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

```
man ipfw
```

Request For Comment Documents

Request for Comments (RFC) documents provide an overview of a protocol or service and details about how the protocol should behave. If you're a novice server administrator, you'll probably find some of the background information in an RFC helpful. If you're an experienced server administrator, you can find all the technical details about a protocol in its RFC document. You can search for RFC documents by number at the website www.faqs.org/rfcs. The "Port Reference" section contains several RFC numbers for various protocols.

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

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

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

Enabling NAT on Mac OS X Server requires detailed control over DHCP, so DHCP is configured separately in Server Admin. To learn more about DHCP, see Chapter 1, “DHCP Service,” on page 7.

Enabling NAT also automatically creates a divert rule to the Firewall configuration.

Starting and Stopping NAT Service

You use Server Admin to start and stop NAT service on your default network interface. Starting NAT service also starts DHCP for the default interface.

To start NAT service:

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

When the service is running, Stop Service becomes available.

Configuring NAT Service

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

To configure NAT service:

- 1 In Server Admin, select NAT from the Computers & Services pane.
- 2 Click Settings.
- 3 Choose the network interface from the “Share your connection from:” pop-up menu.
This interface should be the one that connects to the Internet or external network.
- 4 Click Save.

Monitoring NAT Service

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

Viewing the NAT Status Overview

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

To see the overview:

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

Viewing NAT Activity

When the NAT service is running, it creates a packet divert filter in the IP Firewall service. You can view NAT packet divert events which have been logged by the firewall service. The logs are useful for network troubleshooting and configuration. To troubleshoot NAT, you should create the rule manually and enable logging for the packets allowed by the rule.

To view the NAT divert log:

- 1 In the Terminal application enter:

```
ipfw add 10 divert natd all from any to any via <interface>
```

Where <interface> is the network interface selected in the NAT section of Server Admin.
- 2 In Server Admin, choose Firewall from the Computers & Services list.
- 3 Click Settings.
- 4 Select the Advanced tab.
- 5 Select the rule that was just created.
- 6 Click the Edit button.
- 7 Choose to log packets that match the filter.
- 8 Click OK.
- 9 In Server Admin, choose NAT Service from the Computers & Services list.
- 10 Click Settings.
- 11 Click Logging.
- 12 Enable logging.
- 13 Click Save.
- 14 Click the Log button to view the log.

Where to Find More Information

For more information about natd:

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

```
man natd
```

Request For Comment Documents

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

For NAT descriptions, see RFC 1631 and RFC 3022.

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

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

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

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

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

VPN and Security

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

Authentication Method

Mac OS X Server VPN uses Microsoft's Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2) for authentication. It is also the standard Windows authentication scheme for VPN. This authentication method encodes passwords when they're sent over the network, and stores them in a scrambled form on the server offering good security during network transmission.

This authentication method is the default and available for both transport protocols described in the following section.

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

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

Transport Protocols

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

Point to Point Tunneling Protocol (PPTP)

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

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

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

L2TP/IPSec uses strong IPSec encryption to "tunnel" data to and from the network nodes. It is essentially a combination of Cisco's L2F and PPTP. IPSec requires Security Certificates from a Certificate Authority like Verisign, or a pre-defined shared secret between connecting nodes. The shared secret must be entered on the server as well as a client. It is not a password for authentication, but it is used to generate encryption keys to establish secure tunnels between nodes.

Before You Set Up VPN Service

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

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

If you're using L2TP, you need to have a Security Certificate from a Certificate Authority like Verisign, or a pre-defined shared secret between connecting nodes. If you choose a shared secret, it needs to be secure as well (8-12+ alphanumeric characters with punctuation) and kept secret by the users.

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

Managing VPN Service

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

Starting or Stopping VPN Service

You use Server Admin to start and stop VPN service.

To start or stop VPN service:

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

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

Enabling and Configuring L2TP Transport Protocol

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

To enable L2TP:

- 1 In Server Admin, choose the VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the General tab.
- 4 Select L2TP.
- 5 Enter the shared secret.
- 6 Set the beginning IP address of the allocation range.
- 7 Set the ending IP address of the allocation range.
- 8 Enter the group that has access to VPN login.

You can use the Users & Groups button to browse for a group.

If you leave this blank, all workgroups will have access to VPN login.

- 9 Click Save.

Enabling and Configuring PPTP Transport Protocol

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

To enable PPTP:

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

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

- 6 Set the beginning and IP addresses of the allocation range.
- 7 Enter the group that has access to VPN login.

You can use the Users & Groups button to browse for a group.

If you leave this blank, all workgroups will have access to VPN login.
- 8 Click Save.

Configuring Additional Network Settings for VPN Clients

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

To configure additional network settings:

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

Configuring VPN Network Routing Definitions

Network routing definitions allow you to route data to from some specific address either through the VPN tunnel or the insecure network. For example, you may want all traffic that goes to the LAN IP address range to go through the secure tunnel to the LAN, but make all traffic to other addresses to be routed through the user's normal, unsecured Internet connection. This helps you have a finer control over what goes through the VPN tunnel.

The following definitions are unordered; they apply only the description that most closely matches the packet being routed.

To set routing definitions:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Client Information tab.
- 4 Click the Add button below the routing definition list.
- 5 Enter the address range of the packets to be routed.
- 6 Enter the network mask of the address range to be routed.
- 7 Select the routing destination from the pop-up menu.
Private means to route it through the VPN tunnel.
Public means to use the normal interface with no tunnel.

Monitoring VPN Service

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

Viewing a VPN Status Overview

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

To view the overview:

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

Setting the Log Detail Level for VPN Service

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

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

Non-verbose login is enabled by default.

To set VPN log detail:

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

Setting the VPN Log Archive Interval

Mac OS X Server can automatically archive VPN service logs after a certain amount of time. Each archive log is compressed and uses less disk space than the original log file. You can customize the schedule to archive the logs after a set period of time, measured in days.

To set up the log archive interval:

- 1 In Server Admin, choose VPN Service from the Computers & Services list.
- 2 Click Settings.
- 3 Select the Logging tab.
- 4 Select "Archive Log every ____ days"
- 5 Enter the log archive rollover interval you want.
- 6 Click Save.

Viewing the VPN Log

You'll need to monitor VPN logs to ensure smooth operation of your Virtual Private Network. The VPN logs can help you troubleshoot problems.

To view the log:

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

Viewing VPN Client Connections

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

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

To view client connections:

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

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. See the website www.ietf.org/ids.by.wg/ipsec.html for more information.

Request For Comment Documents

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

- For L2TP description, see RFC 2661.
- For PPTP description, see RFC 2637.

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

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

How NTP Works

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

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

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

Using NTP on Your Network

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

Setting Up NTP Service

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

Additionally, you'll need to make sure your firewall allows NTP queries out to an authoritative time server on UDP port 123, and incoming queries from local clients on the same port. See Chapter 3, "IP Firewall Service," on page 43 for more information on configuring your firewall.

To set up NTP service:

- 1 Make sure your server is configured to "Set Date & Time automatically."
This setting is in the Date & Time pane of System Preferences, or the Server Admin Settings pane for the server.
- 2 Open Server Admin, and select the server you want to act as a time server.
- 3 Click Settings.
- 4 Select the Advanced tab.
- 5 Select Enable NTP.
- 6 Click Save.

Configuring NTP on Clients

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

To configure NTP on clients:

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

Where to Find More Information

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

Request For Comment Documents

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

The official specification of NTP version 3 is RFC 1305.

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

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

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

IPv6 fixes some of these problems and helps others. It adds improvements in areas such as routing and network auto-configuration. It has increased the number of network addresses to over 3×10^{38} , and eliminates the need for NAT. IPv6 is expected to gradually replace IPv4 over a number of years, with the two coexisting during the transition.

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

IPv6 Enabled Services

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

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

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

IPv6 Addresses in the Server Admin

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

IPv6 Addresses

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

Notation

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

IPv6 addresses are generally written in the following form:

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

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

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

or

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

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

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

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

E3C5:4AC8:192.168.100.32

IPv6 Reserved Addresses

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

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

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

IPv6 Addressing Model

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

IPv6 Address Types

IPv6 supports the following three IP address types:

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

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

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

Where to Find More Information

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

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

Request For Comment Documents

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

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

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

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

broadcast The process of transmitting one copy of a stream over the whole network.

byte Eight *bits*.

DHCP (Dynamic Host Configuration Protocol) A protocol used to distribute IP addresses to client computers. Each time a client computer starts up, the protocol looks for a DHCP server and then requests an IP address from the DHCP server it finds. The DHCP server checks for an available IP address and sends it to the client computer along with a *lease period*—the length of time the client computer may use the address.

directory 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.

DoS (denial of service) attack An Internet attack that uses thousands of network pings to prevent the legitimate use of a server.

dynamic IP address An IP address that is assigned for a limited period of time or until the client computer no longer needs the IP address.

filter A “screening” method used to control access to your server. A filter is made up of an IP address and a subnet mask, and sometimes a port number and access type. The IP address and the subnet mask together determine the range of IP addresses to which the filter applies.

firewall Software that protects the network applications running on your server. IP firewall service, which is part of Mac OS X Server software, scans incoming IP packets and rejects or accepts these packets based on a set of filters you create.

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.

HTTP (Hypertext Transfer Protocol) The client/server protocol for the World Wide Web. The HTTP protocol provides a way for a web browser to access a web server and request hypermedia documents created using HTML.

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.

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).

IP (Internet Protocol) Also known as IPv4. A method used with Transmission Control Protocol (TCP) to send data between computers over a local network or the Internet. IP delivers packets of data, while TCP keeps track of data packets.

IP address A unique numeric address that identifies a computer on the Internet.

IP subnet A portion of an *IP* network, which may be a physically independent network segment, that shares a network address with other portions of the network and is identified by a subnet number.

IPSec A security addition to *IP*. A protocol that provides data transmission security for *L2TP VPN* connections. IPSec acts at the network layer, protecting and authenticating IP packets between participating IPSec *nodes*.

IPv6 (Internet Protocol Version 6) The next generation communication protocol to replace IP (also known as IPv4). IPv6 allows a greater number of network addresses and can reduce routing loads across the Internet.

ISP (Internet service provider) A business that sells Internet access and often provides web hosting for ecommerce applications as well as mail services.

L2TP (Layer Two Tunneling Protocol) A network transport protocol used for *VPN* connections. It is essentially a combination of Cisco's L2F and PPTP. L2TP itself is not 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.

load balancing The process of distributing the demands by client computers for network services across multiple servers in order to optimize performance by fully utilizing the capacity of all available servers.

local domain A directory domain that can be accessed only by the computer on which it resides.

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 host The computer that provides your mail service.

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.

MS-CHAPv2 (Microsoft's Challenge Handshake Authentication Protocol version 2) The standard Windows authentication scheme for *VPN*. This authentication method encodes passwords when they are sent over the network and stores them in a scrambled form on the server. It offers good security during network transmission.

multicast An efficient, one-to-many form of streaming. Users can join or leave a multicast but cannot otherwise interact with it.

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 See *DNS (Domain Name System)*.

NAT (Network Address Translation) A method of connecting multiple computers to the Internet (or any other IP network) using one IP address. NAT converts the IP addresses you assign to computers on your private, internal network into one legitimate IP address for Internet communications.

network interface Your computer's hardware connection to some network. This includes (but is not limited to) Ethernet connections, Airport cards, and FireWire connections.

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.

NTP (network time protocol) A network protocol used to synchronize the clocks of computers across a network to some time reference clock. NTP is used to ensure that all the computers on a network are reporting the same time.

Open Directory The Apple directory services architecture, which can access authoritative information about users and network resources from directory domains that use *LDAP*, *NetInfo*, or Active Directory protocols; BSD configuration files; and network services.

open relay A server that receives and automatically forwards mail to another server. Junk mail senders exploit open relay servers to avoid having their own mail servers blacklisted as sources of *spam*.

packet A unit of data information consisting of header, information, error detection, and trailer records. *QTSS* uses *TCP*, *UDP*, and *IP* packets to communicate with streaming clients.

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 or not data packets are allowed to traverse a local network. “Port” usually refers to either a *TCP* or *UDP* port.

protocol A set of rules that determines how data is sent back and forth between two applications.

PTR (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.

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.

Rendezvous A protocol developed by Apple for automatic discovery of computers, devices, and services on IP networks. This proposed Internet standard protocol is sometimes referred to as “ZeroConf” or “multicast DNS.” For more information, visit www.apple.com or www.zeroconf.org.

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.

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.

slave zone The DNS zone records held by a secondary DNS server. A slave zone receives its data by *zone transfers* from the *master zone* on the primary DNS server.

SLP (Service Location Protocol) DA (Directory Agent) A protocol that registers services available on a network and gives users easy access to them. When a service is added to the network, the service uses SLP to register itself on the network. SLP/DA uses a centralized repository for registered network services.

SMTP (Simple Mail Transfer Protocol) A protocol used to send and transfer mail. Its ability to queue incoming messages is limited, so SMTP usually is used only to send mail, and POP or IMAP is used to receive mail.

spam Unsolicited email; junk mail.

SSL (Secure Sockets Layer) An Internet protocol that allows you to send encrypted, authenticated information across the Internet.

static IP address An IP address that is 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 stratum servers are available (2, 3, and so forth); each takes its time from a lower-numbered stratum server.

subnet A grouping on the same network of client computers that are organized by location (different floors of a building, for example) or by usage (all eighth-grade students, for example). The use of subnets simplifies administration.

TCP (Transmission Control Protocol) A method used along with the Internet Protocol (IP) to send data in the form of message units between computers over the Internet. IP takes care of handling the actual delivery of the data, and TCP takes care of keeping track of the individual units of data (called packets) into which a message is divided for efficient routing through the Internet.

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 (text) record A DNS record type that stores a text string for a response to a DNS query.

UCE (unsolicited commercial email) See *spam*.

UDP (User Datagram Protocol) A communications method that uses the Internet Protocol (IP) to send a data unit (called a datagram) from one computer to another in a network. Network applications that have very small data units to exchange may use UDP rather than TCP.

unicast The one-to-one form of streaming. If RTSP is provided, the user can move freely from point to point in an on-demand movie.

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.”

VPN (Virtual Private Network) A network that uses encryption and other technologies to provide secure communications over a public network, typically the Internet. VPNs are generally cheaper than real private networks using private lines but rely on having the same encryption system at both ends. The encryption may be performed by firewall software or by routers.

WAN (wide area network) A network maintained across geographically separated facilities, as opposed to a LAN (local area network) within a facility. Your WAN interface is usually the one connected to the Internet.

wildcard A range of possible values for any segment of an IP address.

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.

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

A

AirPort Base Stations
DHCP service and 9

B

BIND 17, 18, 19, 37–40
about 37
configuration File 38
configuring 37–40
defined 37
example 38–40
load distribution 36
zone data files 38

C

CIDR netmask notation 45, 47

D

DHCP servers 8, 40
interactions 9
network location 8
DHCP service 7–16
AirPort Base Stations 9
changing subnets 11
deleting subnets 12
described 7
disabling subnets 14
DNS options 12
DNS Server for DHCP Clients 12
LDAP auto-configuration 9
LDAP options for subnets 13
logs 15
logs for 10
managing 10–14
more information 16
preparing for setup 7–9
setting up 9–10
starting and stopping 10
subnet IP addresses lease times, changing 12
subnet IP address lease times, changing 12
subnets 8
subnets, creating 10

subnet settings 11
uses for 7
viewing client lists 15
viewing leases, client list 15
WINS options for subnets 14
DNS service 17–41
configuring BIND 37–40
described 17
dynamic IP addresses 40–41
load distribution 36
managing 21–30
more information 41
options for DHCP subnets 12
planning 18
preparing for setup 18
servers 18
setting up 18
setup overview 18–20
starting 21
stopping 21
strategies 18–20
usage statistics 29
uses for 17
with mail service 33
domain names
registering 18, 19
DoS (Denial of Service) attacks
preventing 59
dynamic DNS 40–41
Dynamic Host Configuration Protocol
See DHCP
dynamic IP addresses 8

F

filters
editing 54
examples 57–58
filters, IP
adding 48
described 45

H

help 6

- I
 - IANA registration 18
- In 6
- Internet Gateway Multicast Protocol *See* IGMP
- Internet Protocol Version 6 *See* IPv6
- IP addresses
 - assigning 9
 - DHCP and 7
 - DHCP lease times, changing 12
 - dynamic 8
 - dynamic allocation 8
 - IPv6 notation 84
 - leasing with DHCP 7
 - multiple 47
 - precedence in filters 47
 - ranges 47
 - reserved 9
 - static 8
- IP Filter module 61–63
- IP filter rules 61
- IP Firewall
 - starting and stopping 14
- IP Firewall service 43–44
 - about 43
 - adding filters 48
 - Any Port filter 54
 - background 45
 - benefits 44
 - configuring 49–58
 - creating filters 51
 - default filter 54
 - described 43
 - editing filters 54
 - example filters 57–58
 - filters 45–47
 - IP filter rules 61–63
 - logs, setting up 55–56
 - managing 49–59
 - more information 66
 - multiple IP addresses 47
 - NAT packet divert 68
 - planning 48
 - port reference 63–65
 - preparing for setup 45–47
 - preventing Denial of Service (DoS) attacks 59
 - setting up 48–49
 - starting, stopping 49
 - uses for 44
 - viewing logs 55
- ipfw command 61–63
- IPv6
 - addressing 84–85
 - address notation 84
 - available services 84
 - in Server Admin 84

more information 86

L

- load distribution 36
- logging items
 - DHCP activity 10
- logs
 - DHCP 15
 - DNS service 28
 - IP Firewall service 55–56

M

- Mac OS X Server
 - ports used by 63–65
 - setting up 6
- Mac OS X Server Getting Started 6
- Mac OS X systems 63–65
- mail
 - redirecting 33
- Mail Exchange. *See* MX
- mail exchangers 33
- mail servers 33
- mail service
 - using DNS service with 33
- MX (Mail Exchange) records 20, 33
- MX hosts 33

N

- named.conf file 38
- name servers 18
- NAT
 - about 67
 - activity monitor 68
 - configuring 68
 - monitoring 68
 - more information 69
 - packet divert 68
 - starting, stopping 67
 - status overview 68
 - troubleshooting 68
- NetBoot
 - viewing client lists 15
- networks
 - private 36–37
 - TCP/IP networks 36–37
- NTP
 - about 79
 - configuring clients 81
 - more information 81
 - setting up 80
 - time system 79

O

- online help 6

P

ports

- Mac OS X computers 63–65

- TCP ports 63–64

- UDP ports 65

R

- round robin 36

- rules, IP filter 61–63

S

- Server 10, 15, 57, 58, 69

servers

- DHCP servers 40

- name servers 18

- static IP addresses 8

- Stratum time servers 79

- subnet masks 45

- subnets 8

- creating 8, 10

T

TCP/IP

- private networks 36–37

- TCP ports 63–65

- Terminal application 62

- time servers

- Stratum 79

U

- UDP ports 65

- Universal Time Coordinated (UTC) 79

- User Datagram Protocol See UDP

V

VPN

- client connections 77

- logging 76

- routing definitions 75

- viewing logs 77

- viewing status 76