

Securing Your Mac OS X Server

CONTENT

Addressing Physical Security
Setting Service Access
Configuring the Firewall
Virtual Private Networking
Software Update



Securing Your Mac OS X Server

Schoun Regan



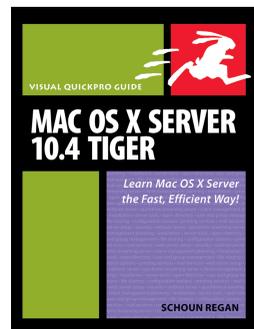
Peachpit Press
1249 Eighth Street
Berkeley, CA 94710
510/524-2178
510/524-2221 (fax)

Find us on the Web at www.peachpit.com

To report errors, please send a note to errata@peachpit.com

Peachpit Press is a division of Pearson Education

This publication originally published as *Mac OS X Server 10.4 Tiger: Visual QuickPro Guide* (ISBN 0-321-36244-6) by Schoun Regan. Copyright © 2006. Published by Peachpit Press.



Notice of Right

All rights reserved. No part of this publication may be reproduced or transmitted in any form by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. For information on getting permission for reprints and excerpts, contact permissions@peachpit.com.

Notice of Liability

The information in this publication is distributed on an “As Is” basis without warranty. While every precaution has been taken in the preparation of the publication, neither the author nor Peachpit Press shall have any liability to any person or entity with respect to any loss or damage caused or alleged to be caused directly or indirectly by the instructions contained in this publication or by the computer software and hardware products described in it.

Trademarks

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this publication, and Peachpit Press was aware of a trademark claim, the designations appear as requested by the owner of the trademark. All other product names and services identified throughout this publication are used in editorial fashion only and for the benefit of such companies with no intention of infringement of the trademark. No such use, or the use of any trade name, is intended to convey endorsement or other affiliation with this publication.

ISBN 0-321-50869-6

Published in the United States of America

A request from the people who worked hard to create this material:

Please respect our copyright. We hope you like what you'll read and learn here, but we also hope you don't pass this material on to others. Thank you very much for your consideration and respect.

—The folks at Peachpit

SECURITY

Planning and implementing security is a necessary part of every system administrator's job. Several aspects are involved, among them physical security, password security, firewalls, encryption, virtual private networks, software updates, and log files. Physical security is much the same as it has been since the lock and key were invented, and password security entails using better methods of storing and using passwords. Everything else relates to being connected to a network, because when it comes down to it, a server that isn't connected to a network isn't a server. And while connecting your server to the Internet allows you to provide services reaching the world over, it simultaneously gives miscreants all over the world access to your server. This brings us to two basic security tenets of running a server: Don't turn on services you don't need and don't serve a wider population than you must.

Apple has been very security-conscious with Mac OS X Server. The out-of-the-box configuration has almost no services turned on by default, allowing you to configure your security first and then turn on the services that you need. You'll need some basic services such as SSH over the local subnet to configure your server and remote access to server tools and directory changes. With Mac OS X Server 10.4, Apple provides comprehensive tools to make managing and monitoring the security of your server easier. The underlying services in Mac OS X Server, including Apache, Samba, OpenLDAP, Postfix, and Jabber, are robust, secure, and actively maintained by the open source community and Apple. This combination of tools and services will help you keep your server running securely.

Addressing Physical Security

Keeping your servers physically secure is one of the first aspects of security that you should consider. If a computer can be physically accessed, then with sufficient time it can be compromised. Because servers are shared resources, their security is usually more important because their compromise would more severely affect your organization. There isn't much point in looking after the other aspects of security until you've addressed physical security.

Preventing booting from various other devices

If your server is in an area where others have physical access to it, they could potentially circumvent the login procedure by power-cycling the computer and restarting in a vulnerable mode. For instance, booting from a system CD, a system DVD, a FireWire hard drive, or a NetBoot drive, the perpetrator will have access to all files on the computer. Booting into single-user mode, the perpetrator will have root access to the computer and all items on that computer.

To enable an Open Firmware lock:

1. Launch the Open Firmware Password application, located in /Applications/Utilities/ on the Mac OS X 10.4 installation disk (**Figure 1.1**).

For Mac OS X 10.1–10.3.9, you can download the Open Firmware Password application from www.apple.com/support/downloads/openfirmwarepassword.html.



Figure 1.1 Recognizing the Open Firmware Password application icon.



Figure 1.2 Launching the Open Firmware Password application brings up this dialog.



Figure 1.3 Setting the password for the firmware using Open Firmware Password.

2. In the dialog that appears, click Change (**Figure 1.2**).
3. Click the “Require password to change Open Firmware settings” check box, enter and verify your password, and click OK (**Figure 1.3**).
4. Enter an administrator name and password in the standard authentication dialog and click OK.
5. Quit Open Firmware Password by choosing Quit from the Open Firmware Password menu.

Startup Keyboard Shortcuts

Apple offers the following keyboard shortcuts, which can be executed after the initial startup chime on PowerPCs:

- ◆ Pressing Option + Command + Shift + Delete during startup attempts to start from a disk other than the primary startup disk.
- ◆ Pressing C during startup attempts to start from a CD or DVD.
- ◆ Pressing N during startup attempts to start from a network server.
- ◆ Pressing T during startup attempts to start in target disk mode.
- ◆ Pressing Shift at boot (after chime) attempts to start in safe boot mode.
- ◆ Pressing V during startup attempts to start in verbose mode.
- ◆ Pressing S during startup attempts to start in single-user mode.

Determining rack and room security

If your server is being used for anything other than personal testing, you must limit physical access to it. Ideally, you have a server room with a locking door to which only a small number of trusted people have keys. If you have a shared server room, you can add security by putting your server in a rack cabinet that locks. If you work in an office that is reasonably secure, a locking rack cabinet may be sufficient, but be sure to review all of the people who have access to the space where your server is, and consider the consequences of a compromised server. What if someone breaks into your building? How critical and/or sensitive is the data on your server? These are questions that should be evaluated in deciding the physical location of your server.

Open Firmware Password Workaround

Locking Open Firmware is a reasonable deterrent, but it does not encrypt the hard drive(s). This strategy can be defeated by changing the amount of RAM in the computer and zapping the Parameter RAM (PRAM). That's a good thing to know in case you ever lock yourself out. In combination with a case lock (to prevent access to changing the amount of RAM), locking the Open Firmware deters the rapid compromise of a machine. It is a best practice to do this in public access areas, open offices, or wherever unauthorized users may have access to the computer.

Power Macs and recent iMacs can be secured from case intrusion with padlocks or locking cables. Older iMacs, PowerBooks, and iBooks only require a screwdriver to change the amount of RAM, defeat an Open Firmware lock, and gain boot variability to the computer. Though the Xserve has a case lock, it uses a hex key and is primarily intended to prevent accidental intrusions. Xserves are most vulnerable to physical intrusion precisely because they are expected to be protected in a secured rack and/or room.

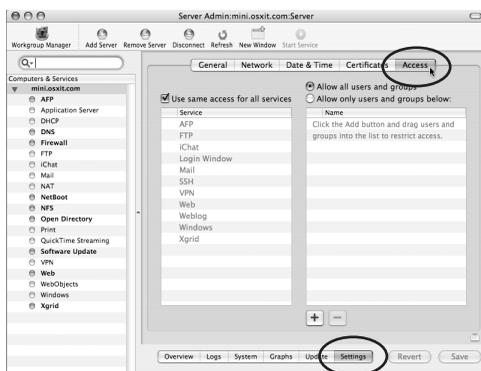


Figure 1.4 Using Server Admin to control access to services.

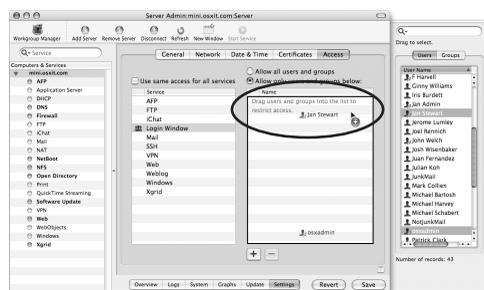


Figure 1.5 Dragging in users from the user list to restrict access to the Login window.

Setting Service Access

Service Access Control Lists (ACLs) are new to Mac OS X Server 10.4. Like the firewall service discussed in the following section, service ACLs are a very powerful security tool—be *very* careful not to lock yourself out of your server! Service ACLs determine which users or groups have access to the services provided by your server. In previous versions of Mac OS X Server, if a service was enabled for one user, it was enabled for all. Service ACLs are an additional security measure that will limit service usage to only those users you define.

The following task shows you how to restrict people from logging into the server, even if they have access to a keyboard and monitor attached to the server.

To restrict access to the Login window:

1. Launch Server Admin and select your server from the Computers & Services list. You don't need to authenticate if you have already added your server to the keychain. Leave Server Admin running for the next several exercises.
2. Click the Settings button and then click the Access tab (**Figure 1.4**).
3. Deselect the “Use same access for all services” check box but select the “Allow only user and groups below” check box.
4. Select Login Window from the Service list below and click the plus button to open the Users and Groups drawer on the right side of the window.
5. Click-and-drag your username and any other users from the drawer to the Name list (**Figure 1.5**).
6. Click the Save button to save your changes and permit only the selected users to log in via the Login window.
7. Test access by attempting to log in as a user not included in the access list.

Configuring the Firewall

A firewall offers protection by selectively filtering network packets: you can prevent unwanted inbound traffic from the Internet, and you can restrict outbound traffic from your internal network. The firewall software in Mac OS X Server, IPFirewall (IPFW), is from the FreeBSD organization and is a very powerful and sophisticated tool that offers protection over both IP version 4 and IP version 6.

Fortunately, Apple has made this much more approachable through the Server Admin tool. The best way to set up your firewall is to start with Mac OS X Server's default settings and then create openings only for the network traffic needed by defining a group of IP addresses and selecting the services you want them to be able to access.

What Are the Default Firewall Settings?

Mac OS X Server 10.4 has default firewall settings that will allow you to configure and run your server over a network. Whether you have an Xserve with no monitor or you are configuring a server thousands of miles away, you will need to allow at least some traffic through your firewall. The default settings include special settings, which have a lock icon next to them and cannot be changed, and minimal management settings, which can be changed if you desire. The default management settings are:

- ◆ SSH—Secure Shell
- ◆ Server Admin SSL, also Web-ASIP
- ◆ Remote Directory Access
- ◆ Serial number support

Serial number support allows traffic to check the legitimacy of your Mac OS X Server 10.4 license; if you deselect the check box for this service, it will turn itself back on. The remaining default services are allowed for any IP address, and may be necessary in order to set up a Mac OS X Server from across the country or from a different subnet. If you will not need such wide access to manage your servers, limit traffic to only the range of IP addresses from which you will be managing them.

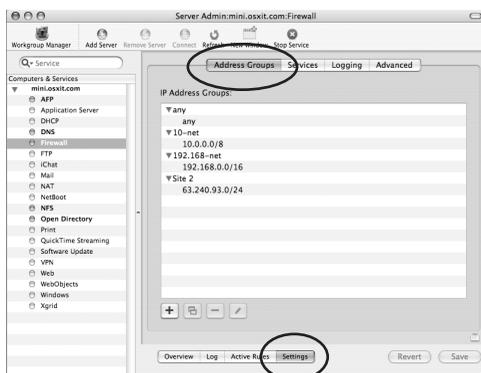


Figure 1.6 Selecting the Firewall group settings in Server Admin.

Using the Command Line with the Firewall

IPFW can be configured from the command line by using the `serveradmin` command or traditional `ipfw` commands. The `serveradmin` commands essentially give you command-line access to the Server Admin way of doing things, and the `ipfw` commands give you access to the traditional, FreeBSD way. The `serveradmin` commands are useful for tasks such as starting or stopping the firewall service, or checking its status. The `ipfw` commands give you a convenient way to add or modify firewall rules.

It is important to understand how the GUI way and the FreeBSD way interact: Looking in the `/private/etc/ipfilter` directory, you'll find `ipfw.conf`, `ipfw.conf.apple`, and `ipfw.conf.default`. Viewing the contents of the `ipfw.conf` file will explain the interaction between the GUI elements and the FreeBSD elements.

Be *very* careful when configuring the firewall to not lock yourself out! Make a habit of confirming that you have left the following management services accessible before you save a firewall configuration:

- ◆ SSH—Secure Shell
- ◆ Server Admin SSL, also Web-ASIP
- ◆ Remote Directory Access

It is a good idea to set up a small shell script to flush all the firewall rules every 15 minutes while you test and work on your firewall. This way, if you do lock yourself out, you only need wait until the 15 minutes rolls around and then you can start again.

When you define address groups, you are telling the firewall to adhere to the rules for just that IP range. For servers with only one Ethernet connection and one IP address, you can delete all other groups except the All group; then, no matter what your computer's IP address, you will be managing the firewall for any IP address that the server uses.

To define address groups:

1. In Server Admin, select the Firewall service for your server in the Computers & Services list.
2. Click the Settings button and then click the Address Groups tab (**Figure 1.6**).

continues on next page

3. Click the plus button to open a dialog where you can provide a group name and add addresses; then click OK to return to the main window (**Figure 1.7**).
4. Confirm that the Address range displayed in the lower part of the IP Address Groups configuration drawer is as you intend it to be, and click OK (Figure 10.6). Your new address group will appear in the IP Address Groups window.
5. When you've finished making changes, click Save.

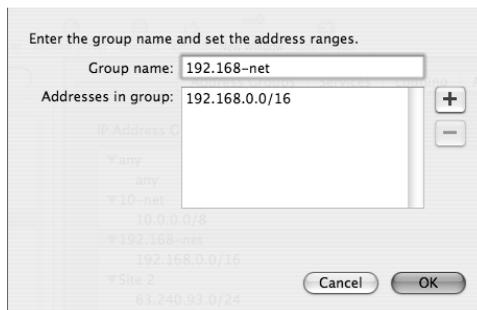


Figure 1.7 Adding a new group to the Firewall group list.

Figuring Out Subnet Mask Notation

There are three ways you can add IP addresses to a firewall address group. The first is to simply enter a single IP address. But what if you wanted to add more than a single IP address at a time—for example, your company's whole address range? In that case you'd need to define a subnet mask, either using Classless Inter-Domain Routing (CIDR) notation or netmask notation. If you haven't spent a lot of time configuring network settings, this may seem like a daunting task. Consult one of many Internet sites devoted to understanding networking and CIDR notation.

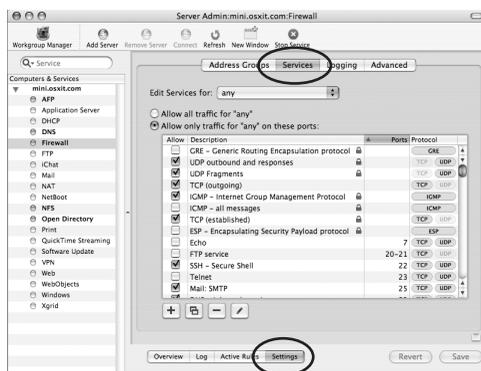


Figure 1.8 Selecting which services are allowed to be accessed through the firewall over the group *any*.



Figure 1.9 Choosing another group over which firewall rules will be applied.

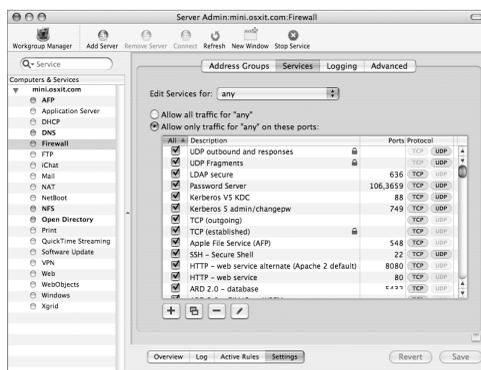


Figure 1.10 Sorting rules by whether they are checked or not.

Allowing access to services

Once you have your IP address groups defined, you can allow services to be accessed by those groups. In the Services pane, you simply select the services you want.

To allow access to services:

1. Select the Firewall service, click the Settings button, and then click the Services tab (**Figure 1.8**).
2. The services with the padlock icon cannot be edited or removed.
3. From the “Edit Services for” pop-up menu, select the address group to which you want to add a service (**Figure 1.9**).
4. Select the “Allow only traffic for ‘any’” on these ports” radio button, and select the service(s) you want to allow below.
5. Click the All column to display selected services at the top of the list (**Figure 1.10**).
6. When you’ve finished making changes, click the Save button and make sure the firewall is running.

Your firewall should now allow the traffic for the service(s) you selected for the IP address group you specified.

✓ Tips

- By repeating the process for various other groups (depending on how many interfaces are active on your server), you can provide the services you want to the computers that need them.
- The best way to implement the firewall is to turn off all possible services except the management tools and ssh and then turn on services as necessary. For example, when using Mac OS X Server as a KDC (Open Directory Master), you may not need to open all the services surrounding Kerberos to make secure connections work.

To add specific firewall rules:

1. Select the firewall service, click the Settings button, and then click the Services tab.
2. From the “Edit Services for” pop-up menu, select the address group for which you want to enable your service.
3. Click the plus button and a new service dialog appears (**Figure 1.11**).

4. Provide a service name, port number, and then click OK (**Figure 1.12**).

Your new service is added to the services list for all of the service groups, but you must select the check box to make that service accessible through the firewall (**Figure 1.13**).

5. When you’ve finished making changes, click Save.

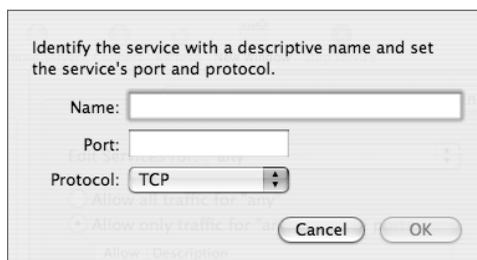


Figure 1.11 Clicking the plus button allows a new service to be listed in the service list.

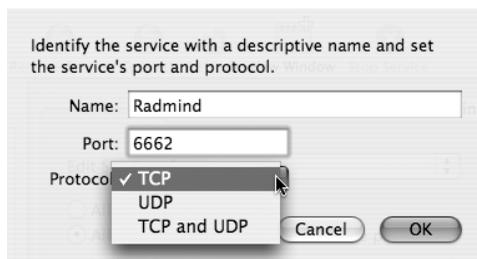


Figure 1.12 Entering data into the new service dialog.

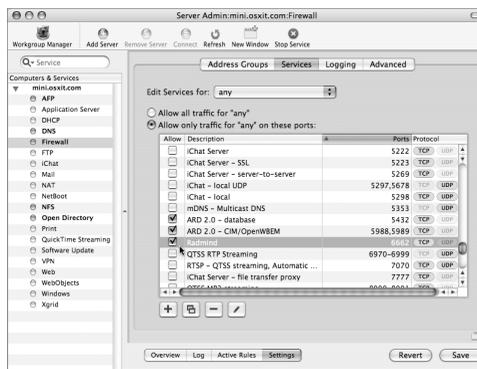


Figure 1.13 Selecting the newly added service in the list.

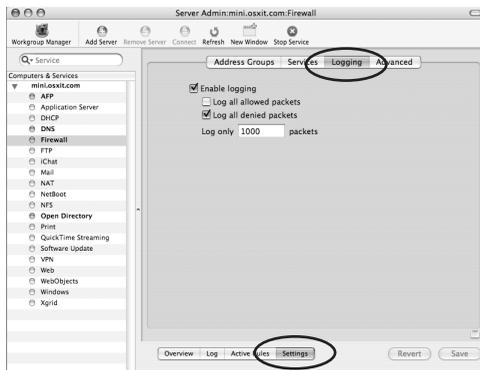


Figure 1.14 Enabling firewall logging via the Logging tab.

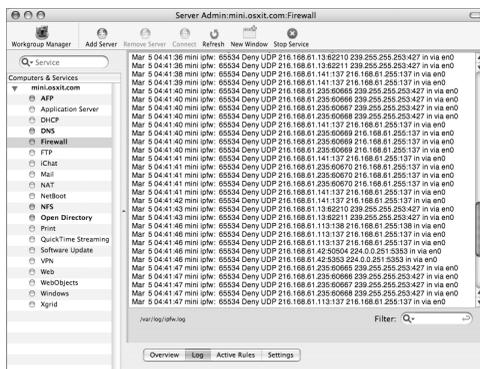


Figure 1.15 Viewing the firewall log from within Server Admin.



Figure 1.16 Restricting the log view using the filter.

Setting logging options for the firewall

The logging options for the firewall in Mac OS X Server are quite simple. Generally, you will log all allowed and/or denied packets for analysis or troubleshooting rather than leaving those options selected all the time. Logging every packet can result in some pretty big log files, so be careful! You can limit how many packets are logged, too.

To set logging options:

1. Select the Firewall service, click the Settings button, and then click the Logging tab (**Figure 1.14**).
2. Click the “Log all denied packets” check box and set the value to “Log only 1000 packets” for a relatively small log file.
3. To test the limits you just set, attempt to connect to your Mac OS X Server with a service blocked by the firewall from any other computer.
4. Click the Log button at the bottom of the pane to display entries for all denied packets (**Figure 1.15**).

For example, type `telnet xxx.xxx.xxx.xxx 115` (the *x*'s represent the IP address of your server; 115 is the port you want to test).

✓ Tip

- To more easily isolate the packets you are looking for, type a colon (:) followed by the port number of the service you are looking for (**Figure 1.16**).

Configuring advanced settings and rules

Mac OS X Server's firewall also offers two Stealth Mode advanced options that you can use to give your server a less conspicuous presence on the Internet. When you select the Enable for TCP and Enable for UDP check boxes, packets that your server receives on closed ports are simply dropped; this is the same response an attacker would get if there were no computer present (**Figure 1.17**).

Though Mac OS X Server provides firewall rules for most of the services you will need, there may be times when you need to add your own custom rules. One reason you'd want to make an advanced rule is to allow protocols other than Transmission Control Protocol (TCP) or User Datagram Protocol (UDP)—such as Internet Control Message Protocol (ICMP), Internet Group Management Protocol (IGMP), Encapsulating Security Payload (ESP), or the protocol used to encapsulate static packets in an IP header, IPEncap.

Another reason you'd want to make an advanced rule is to allow different traffic on different network interfaces: All of the entries under the Services tab apply to all the network interfaces of your server. So if you have multiple physical or virtual interfaces defined in your Network preference pane, you can create custom rules under the Advanced tab to set up independent rules for each interface.

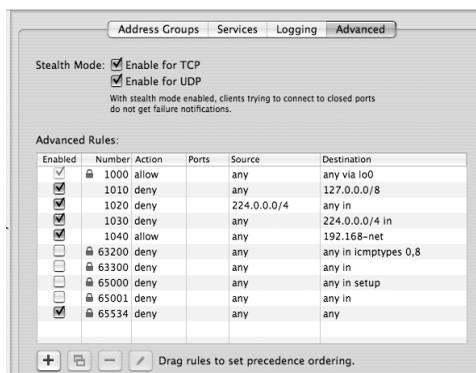


Figure 1.17 The Advanced tab of the firewall settings allows for rules to be moved in order of importance.

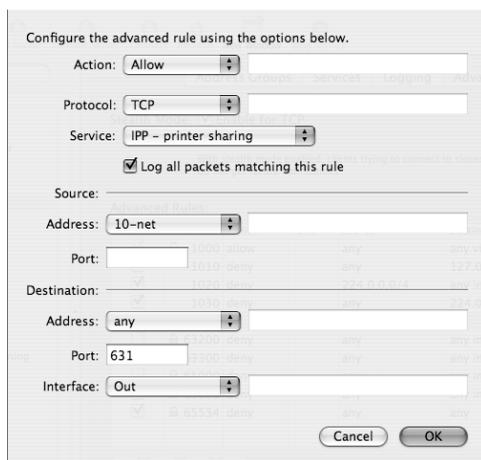


Figure 1.18 Adding a new advanced rule opens this dialog.

Password Security

Password security and physical security share the dubious distinction of often being neglected weak links in a security system. It is frustrating to think how quickly a well-planned security system can be compromised by a poor password choice. With Mac OS X Server, you have the ability to enforce password policies per user or globally for all users.

To create an advanced rule, you must understand what you want to accomplish. Here are some main options to be explored (**Figure 1.18**):

- ◆ From the Action pop-up menu, choose Allow or Deny.
- ◆ From the Protocol pop-up menu, choose the type of protocol affected by the firewall.
- ◆ Choose a specific service from the Service pop-up menu or select Other if the preset service is not in the list.
- ◆ In the Source area, the Address pop-up menu provides you with each IP address group from the Address Group tab. You can also enter specific IP addresses or ranges in the field.
- ◆ In the Port field, you can specify which source port will be used with this specific rule.
- ◆ In the Destination area, the Address pop-up menu provides the same IP address groups as in the Source area, and the port is sometimes preselected based on your service choice.
- ◆ From the Interface pop-up menu, choose whether the traffic is coming into or out of your server.

After you set up your customized rule, click OK and then save your changes. You may want to drag the rule higher or lower in the list depending on where you want the rule to fall with respect to your other rules.

Recovering from a Lockout

If you are in the unenviable position of being locked out of your server, the firewall will need to be reset to its default settings. All is not lost, but you will need to have physical access to the server, and be able to boot it into single-user mode.

1. Disconnect the server from the Internet.
2. Restart in single-user mode by pressing `Cmd+S`.
3. Follow the onscreen prompts by typing `/sbin/fsck -yf`.
4. When the `fsck` file system check has successfully run, mount the file system:
`/sbin/mount -uw /`.
5. Rename the `ipfw` configuration file and the address groups file:

```
cd /private/etc/ipfilter
mv ipfw.conf ipfw.conf.old
mv ip_address_groups.conf ip_address_groups.conf.old
```
6. Flush the firewall rules:
`ipfw -f flush`
7. Edit `/private/etc/hostconfig` to confirm that the `IPFILTER` setting reads as follows:
`IPFILTER=-YES-`

This will ensure that your server starts with the firewall active.
8. Restart your server and check the firewall configuration.
9. Reconnect your server to the Internet.

Virtual Private Networking

Virtual private networks, or VPNs, extend the reach of your security measures considerably by allowing your remote users to keep their communication with your LAN secure. Not only does that help to protect the data flowing in and out of your organization, but you also don't have to compromise the security of your network in order to provide convenient remote access for your users. VPNs create a secure, encrypted tunnel through which your users can connect to your LAN. Any services that aren't already encrypted will be protected by the VPN.

It also makes it easier for your remote users to access certain resources that are restricted to your LAN, such as services restricted by the firewall to IP addresses on the LAN. In short, a VPN allows a remote computer to behave as though it is directly connected to your LAN.

The VPN service gives you the option of adding other routes to your routing table as well. For example, if you have two different networks at your business and each network has different IP address ranges, you can add these networks to your VPN client information. If you are not sure which IP ranges need to be made available, consult the network administrator at your organization.

There are two methods of connecting to a network via a VPN connection using Server Admin, and each of these contains variances as well. The more secure of the two is Layer Two Tunneling Protocol over IP/Sec (L2TP). The other, less secure method is Point-to-Point Tunneling Protocol (PPTP). L2TP has various options (which will be discussed in the following task). PPTP only offers to add lower 40-bit level encryption to the existing 128-bit encryption.

To enable the VPN service:

1. In Server Admin, select the VPN service for your server in the Computers & Services list.
2. Click the Settings button and then select the L2TP tab (Figure 1.19).
3. Click the Enable L2TP check box and fill in starting and ending IP addresses to define the range (Figure 1.20).
4. From the PPP Authentication pop-up menu, choose either Microsoft's implementation of the Challenge Handshake Authentication Protocol, version 2 (MS-CHAPv2) or Kerberos.
5. In the IPsec Authentication area, select either a Shared Secret (that can be seen when typed in) or a Certificate.
6. To configure PPTP, select the PPTP tab and enter the appropriate information for clients connecting over PPTP (Figure 1.21).

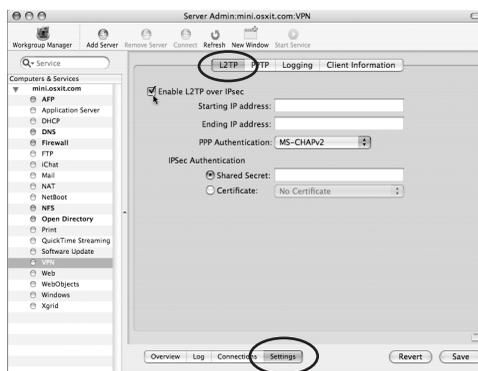


Figure 1.19 Select the check box to enable L2TP over IP/Sec for the VPN Service.

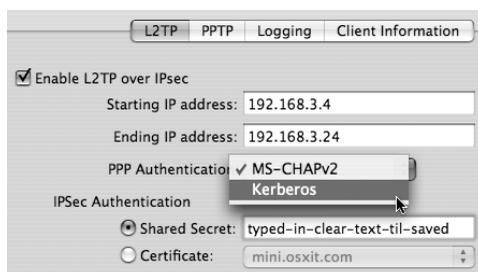


Figure 1.20 Setting various L2TP options.

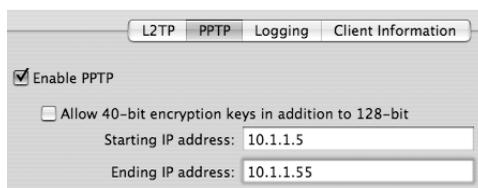


Figure 1.21 PPTP options are minimal.

The screenshot shows a configuration window with four tabs: L2TP, PPTP, Logging, and Client Information. The Client Information tab is active. It contains two text input fields: 'DNS servers:' with the value '192.168.3.2' and 'Search domains:' with the value 'myinternaldomain.com'.

Figure 1.22 Client information can include DNS and domain information and...

The screenshot shows a table titled 'Network Routing Definition:' with three columns: Network Address, Network Mask, and Network Type.

Network Address	Network Mask	Network Type
172.16.9.0	255.255.0.0	Private
216.168.61.0	255.255.255.0	Public

Figure 1.23 ...both public and private routes that will be added to the user's routing table.

7. Click the Client Information tab and fill in the DNS and search domain so the connecting machine can access internal devices accurately (**Figure 1.22**).
8. Add any public or private routes to be listed in the user's routing table (**Figure 1.23**).
9. When you've finished making changes, click the Save button and start the VPN service.

✓ Tip

- If you are troubleshooting a VPN connection within Internet Connect, select Options from the Connect menu and select the "Use Verbose Logging" and "Send all traffic over VPN connection" check boxes to assist you in locating the issue.

Internet Connect application setup

The Internet Connect application also needs to be configured to allow Mac OS X clients to connect to the VPN. When this is done, a new virtual interface is added to the current network location's interface list. If you have more than one location, you should repeat the process that follows for each location, thereby adding the virtual interface to each location.

To configure the Internet Connect application:

1. Launch the Internet Connect application located in /Applications (**Figure 1.24**), and select the VPN icon, if available, in the application's toolbar.

or

Choose New VPN Connection from the File menu (**Figure 1.25**).

2. In the dialog that appears, choose the appropriate connection method and click Continue (**Figure 1.26**).

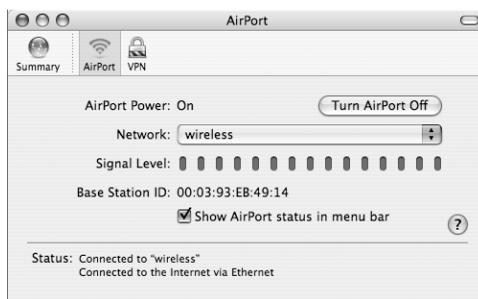


Figure 1.24 Opening the Internet Connect application...



Figure 1.25 ...to add a new VPN connection from the File menu.



Figure 1.26 When adding a new VPN connection, a new VPN dialog appears.



Figure 1.27 Choosing to edit the new VPN configuration.

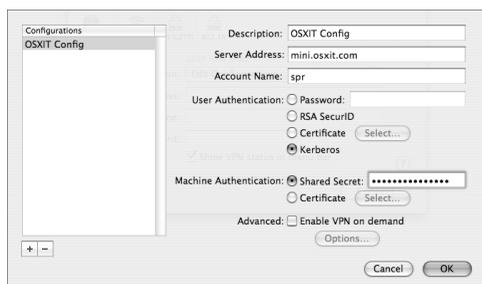


Figure 1.28 Entering various data into the new VPN dialog.

3. In the Internet Connect application, choose Edit Configurations from the Configuration pop-up menu (**Figure 1.27**).
4. In the description field that subsequently appears, provide a name for this VPN connection, in case you have more than one to choose from (**Figure 1.28**).
5. In the Server Address field, provide the IP address or domain name of the VPN server.
6. In the Account Name field, enter the short name of the user who will connect to the VPN.
7. In the User Authentication area, select a method of authentication.
If you are using CryptoCard instead of RSA, contact CryptoCard for information on compatibility at www.cryptocard.com.
8. In the Machine Authentication area, select either a Shared Secret or Certificate.

continues on next page

9. In the Advanced area, you can click the “Enable VPN on demand” check box and then click Options to display a dialog where you can provide a domain name to trigger the connection (**Figure 1.29**).
10. Click OK and then click Connect to connect to your VPN service (**Figure 1.30**).

✓ Tips

- You can select the “Show VPN status in menu bar” check box in the Internet Connect application to connect to the VPN directly from the menu bar.
- There are other, more complex connection methods to other types of VPN servers that are beyond the scope of this book.

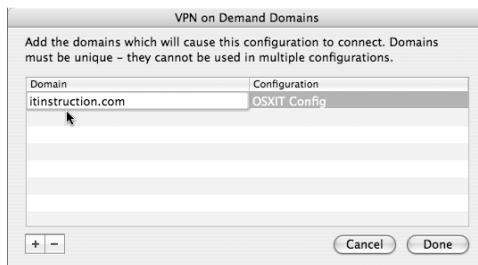


Figure 1.29 Adding a domain to enable VPN on demand.



Figure 1.30 Internet Connect showing VPN connection.

Firewall VPN Ports

The firewall should have the following ports open, depending on your VPN configuration:

- ◆ 500—Internet Security Association and Key Management Protocol/Internet Key Exchange (VPN ISAKMP/IKE)
- ◆ 1701—VPN L2TP—Layer-Two Tunneling Protocol
- ◆ 1723—VPN PPTP—Point-to-Point Tunneling Protocol
- ◆ 4500—IKE NAT Traversal

Always back up your firewall settings with the tear-off configuration plist icon in the lower-right corner of the Settings window so that you can return to a set state, if necessary.

Software Update

Traditional wisdom with software updates is that you wait for other administrators who are clamoring to be first on the bandwagon to install the updates and find out what, if any, issues the updates caused. Once there has been some public testing and it looks as though the update helps more than it hurts, you go ahead and install it.

But now, in this day and age of constant Internet attacks, it is increasingly risky to leave a computer unpatched and vulnerable, especially with respect to security updates. You now have to balance these two approaches carefully. For updates that contain new features and functionality, there is less reason to rush to install, especially on a server. On the other hand, it is probably advisable to install security updates as soon as possible.

Fortunately, because patches are qualified by both Apple and the open source community, security updates are usually timely, reliable, and unlikely to create problems. The Software Update Server can permit users to run the application as they normally would, except for the fact that they are actually connecting to your server.

Another option when you're securing a server is to restrict the users who can access the server via ssh. Since ssh is turned on by default when you install and configure Mac OS X Server, a good password may not be enough to keep others out. You can use open source software contained in Mac OS X Server to generate a key and then give the key to individuals who are permitted to log in via ssh.

To create an ssh key:

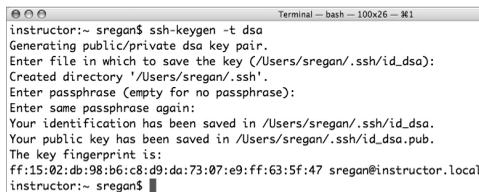
1. Launch the Terminal application located in /Applications/Utilities, enter the command `ssh-keygen -t dsa`, and press Return (**Figure 1.31**).

This command generates the two files necessary for the encryption using the type dsa: the file `id_dsa` contains the private part of the key, and the file `id_dsa.pub` contains the public part of the key.

2. In the line that appears after pressing Return from step 1, you do not need to provide a filename. So you can press Return again.
3. On the line that now appears, you can add an optional passphrase to encrypt the private key with Triple-DES encryption for even more security.
4. Copy the `id_dsa.pub` file to your administrator's home folder on the server. If it doesn't already exist, create a folder called `.ssh` in the home folder on the server, move the `id_dsa.pub` file into the `.ssh/` folder, and rename it `authorized_keys` (**Figure 1.32**).

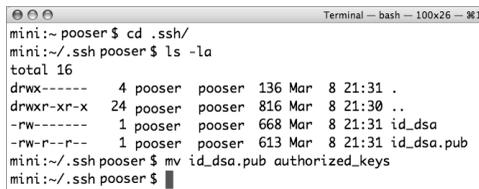
5. On your server, open and edit the `/etc/sshd_config` file using any command-line editor by changing the following lines from:

```
#PasswordAuthentication yes
#ChallengeResponseAuthentication yes
to
PasswordAuthentication no
ChallengeResponseAuthentication no
```



```
Terminal — bash — 100x26 — #11
instructor:~ sregan$ ssh-keygen -t dsa
Generating public/private dsa key pair.
Enter file in which to save the key (/Users/sregan/.ssh/id_dsa):
Created directory '/Users/sregan/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /Users/sregan/.ssh/id_dsa.
Your public key has been saved in /Users/sregan/.ssh/id_dsa.pub.
The key fingerprint is:
ff:15:02:db:98:b6:c8:d9:da:73:07:e9:ff:63:5f:47 sregan@instructor.local
instructor:~ sregan$
```

Figure 1.31 Using the Terminal to create a public/private key pair.



```
Terminal — bash — 100x26 — #11
mini:~ pooser $ cd .ssh/
mini:~/.ssh pooser $ ls -la
total 16
drwx----- 4 pooser pooser 136 Mar  8 21:31 .
drwxr-xr-x 24 pooser pooser 816 Mar  8 21:30 ..
-rw----- 1 pooser pooser 668 Mar  8 21:31 id_dsa
-rw-r--r-- 1 pooser pooser 613 Mar  8 21:31 id_dsa.pub
mini:~/.ssh pooser $ mv id_dsa.pub authorized_keys
mini:~/.ssh pooser $
```

Figure 1.32 Renaming the file on the server.

6. You may want to restart your server at this point because attempting to stop and restart ssh can cause you to lock yourself out of your server.
7. Attempt to log in again from the Mac OS X computer that generated the keys and enter the passphrase when requested.
Any hackers who attempt to log in using your short name will not be able to log in without the private key on their computer.

✓ Tips

- It's a good idea to also copy the .ssh directory on your server from your administrator's home folder to root's home folder, thus keeping a direct remote root login safer by restricting it to key access.
- Make a backup copy of the keys, and install them on all machines that may require other administrators to ssh into the server. You may also wish to copy the authorized_keys file from one administrator account to other accounts or to generate individual keys for each user.
- If you don't enter a passphrase when you generate a key, then when a user opens the Terminal and attempts to log in with the admin account, they won't be asked for a password or passphrase. This can be dangerous, because the client computer can access the server by opening the Terminal and typing in `ssh admin_name@ip_address`.