



Mac OS X Server

High Availability Administration
For Version 10.4 or Later

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

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

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

Apple
1 Infinite Loop
Cupertino CA 95014-2084
www.apple.com

The Apple logo is a trademark of Apple Computer, Inc., registered in the U.S. and other countries. Use of the “keyboard” Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

Apple, the Apple logo, AppleShare, AppleTalk, Mac, Macintosh, QuickTime, Xgrid, and Xserve are trademarks of Apple Computer, Inc., registered in the U.S. and other countries. Finder is a trademark of Apple Computer, Inc.

Adobe and PostScript are trademarks of Adobe Systems Incorporated.

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

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

019-0175/03-24-05

Contents

Preface	5 About This Guide
	5 What's New in Version 10.4
	5 What's in This Guide
	6 Using Onscreen Help
	6 The Mac OS X Server Suite
	8 Getting Documentation Updates
	8 Getting Additional Information
Chapter 1	9 Overview of High Availability
	9 What is High Availability?
	9 It's a Big Field
	10 High Availability Software Solutions
	10 High Availability Hardware Solutions
Chapter 2	11 Configuring IP Failover
	12 IP Failover Overview
	14 Acquiring Master Address—Chain of Events
	15 Releasing Master Address—Chain of Events
	15 IP Failover Setup
	16 Connecting the Master and Backup Servers to the Same Network
	16 Connecting the Master and Backup Servers Together
	17 Configuring the Master Server for Failover
	17 Configuring the Backup Server for Failover
	18 Viewing the IP Failover Log
Chapter 3	19 Configuring Other Aspects of High Availability
	19 Eliminating Single Points of Failure
	20 Using Xserve and Xserve RAID
	20 Using Xserve
	20 Using Xserve RAID
	21 Using Backup Power
	21 Using UPS with Xserve RAID
	22 Using UPS With Xserve
	23 Setting Up Your Server for Automatic Reboot

	24	Ensuring Proper Operational Conditions
	24	Open Directory Replication
	24	Protecting Server Security
	25	Link Aggregation
	26	The Link Aggregation Control Protocol (LACP)
	27	Link Aggregation Scenarios
	29	Setting Up Link Aggregation in Mac OS X Server
	31	Monitoring Link Aggregation Status
	32	Load Balancing
	33	Backup
	33	Developing a Backup Strategy
	35	Command-Line Backup and Restoration Tools
Chapter 4	37	Monitoring Server Availability
	37	Console
	38	Server Monitor
	38	RAID Admin
	39	Disk Monitoring Tools
	40	Network Monitoring Tools
Chapter 5	41	Solving IP Failover Problems
	41	Service Doesn't Automatically fail over When Primary Service Is Unavailable
	41	Email Recipient Isn't Getting Notifications
	41	Failover Took Place, But Still Having Problems
Glossary	43	
Index	51	

About This Guide

This guide describes how to ensure high availability of your Mac OS X Server services.

This guide describes how you can use IP failover, link aggregation, load balancing, and other software and hardware configurations to ensure high availability of your Mac OS X Server services.

What's New in Version 10.4

Mac OS X Server version 10.4 offers major enhancements that support high availability:

- Link aggregation—Mac OS X Server now supports link aggregation, which you configure in the Network pane of System Preferences.
- New daemon manager—The `launchd` daemon has replaced the `watchdog` process. Like `watchdog`, `launchd` ensures that processes that are supposed to be running are always running.

What's in This Guide

This guide includes the following chapters:

- Chapter 1, “Overview of High Availability,” explains the concept of high availability.
- Chapter 2, “Configuring IP Failover,” describes how you can set up IP failover in Mac OS X Server.
- Chapter 3, “Configuring Other Aspects of High Availability,” describes how you can use link aggregation, load balancing, and other software and hardware configurations to ensure high availability.
- Chapter 4, “Monitoring Server Availability,” describes the different tools that you can use to monitor the availability of your services.
- Chapter 5, “Solving IP Failover Problems,” describes common problems and provides information on how to solve them.
- The glossary provides brief definitions of terms used in this guide.

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

Using Onscreen Help

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

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

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

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

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

The Mac OS X Server Suite

The Mac OS X Server documentation includes a suite of guides that explain the services and provide instructions for configuring, managing, and troubleshooting the services. All of the guides are available in PDF format from:

www.apple.com/server/documentation/

This guide ...	tells you how to:
<i>Mac OS X Server Getting Started for Version 10.4 or Later</i>	Install Mac OS X Server and set it up for the first time.
<i>Mac OS X Server Upgrading and Migrating to Version 10.4 or Later</i>	Use data and service settings that are currently being used on earlier versions of the server.
<i>Mac OS X Server User Management for Version 10.4 or Later</i>	Create and manage users, groups, and computer lists. Set up managed preferences for Mac OS X clients.
<i>Mac OS X Server File Services Administration for Version 10.4 or Later</i>	Share selected server volumes or folders among server clients using these protocols: AFP, NFS, FTP, and SMB/CIFS.

This guide ...	tells you how to:
<i>Mac OS X Server Print Service Administration for Version 10.4 or Later</i>	Host shared printers and manage their associated queues and print jobs.
<i>Mac OS X Server System Image and Software Update Administration for Version 10.4 or Later</i>	Use NetBoot and Network Install to create disk images from which Macintosh computers can start up over the network. Set up a software update server for updating client computers over the network.
<i>Mac OS X Server Mail Service Administration for Version 10.4 or Later</i>	Set up, configure, and administer mail services on the server.
<i>Mac OS X Server Web Technologies Administration for Version 10.4 or Later</i>	Set up and manage a web server, including WebDAV, WebMail, and web modules.
<i>Mac OS X Server Network Services Administration for Version 10.4 or Later</i>	Set up, configure, and administer DHCP, DNS, VPN, NTP, IP firewall, and NAT services on the server.
<i>Mac OS X Server Open Directory Administration for Version 10.4 or Later</i>	Manage directory and authentication services.
<i>Mac OS X Server QuickTime Streaming Server Administration for Version 10.4 or Later</i>	Set up and manage QuickTime streaming services.
<i>Mac OS X Server Windows Services Administration for Version 10.4 or Later</i>	Set up and manage services including PDC, BDC, file, and print for Windows computer users.
<i>Mac OS X Server Migrating from Windows NT for Version 10.4 or Later</i>	Move accounts, shared folders, and services from Windows NT servers to Mac OS X Server.
<i>Mac OS X Server Java Application Server Administration For Version 10.4 or Later</i>	Configure and administer a JBoss application server on Mac OS X Server.
<i>Mac OS X Server Command-Line Administration for Version 10.4 or Later</i>	Use commands and configuration files to perform server administration tasks in a UNIX command shell.
<i>Mac OS X Server Collaboration Services Administration for Version 10.4 or Later</i>	Set up and manage weblog, chat, and other services that facilitate interactions among users.
<i>Mac OS X Server High Availability Administration for Version 10.4 or Later</i>	Manage IP failover, link aggregation, load balancing, and other hardware and software configurations to ensure high availability of Mac OS X Server services.

This guide ...	tells you how to:
<i>Mac OS X Server Xgrid Administration for Version 10.4 or Later</i>	Manage computational Xserve clusters using the Xgrid application.
<i>Mac OS X Server Glossary: Includes Terminology for Mac OS X Server, Xserve, Xserve RAID, and Xsan</i>	Interpret terms used for server and storage products.

Getting Documentation Updates

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

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

Getting Additional Information

For more information, consult these resources:

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

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

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

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

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

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

This chapter describes the concept of high availability and what you need to do to apply it to your Mac OS X Server deployments.

Because of service level guarantees, business requirements, or established industry regulations, many companies require the highest possible availability of computing services. Today, events that were previously viewed as minor, unplanned outages can severely impair business operations. In response, Apple's Mac OS X Server products are designed to ensure high availability and manage system outages.

What is High Availability?

In the context of network services, high availability refers to almost continuous availability of services. In many organizations, a highly available system or service must be usable 99% of the time. Other organizations require even greater availability: 99.99% or 99.999% in some cases.

The following table puts these numbers in perspective:

Availability	Uptime (approx.)	Downtime (approx.)
99%	361.8 days	84 hours
99.9%	364.9 days	8.5 hours
99.999%	365.3 days	5 minutes

Once you have determined the level of availability that your services must provide, you can calculate the allowable down time and use the resulting number to determine how you'll configure your hardware and software to meet that level of availability.

It's a Big Field

High availability is a big field. As you'll learn in this guide, it takes a combination of software and hardware solutions to ensure high availability.

High Availability Software Solutions

Mac OS X Server provides several features to ensure the highest possible availability:

- IP failover service allows you to provide redundancy. When the master server fails, a backup server automatically assumes the role of master server with minimal or no disruption to users.
- The `launchd` daemon makes sure that processes specified in its configuration file are always running. If one of these processes quits, `launchd` restarts it. For more information about `launchd`, refer to the command-line administration guide.
- Open Directory service allows you to replicate your directory and authentication services so that users see little disruption in service if a directory system fails or becomes unreachable.
- The Energy Saver pane of System Preferences lets you set an option for restarting the computer in case it goes down.
- Monitoring utilities such as Server Monitor and RAID Admin notify you via email when the operating conditions for any component exceed your predefined thresholds, so you can respond quickly to prevent or repair the problem.
- The `rsync` command along with other third-party tools allow you to mirror your data so that in the event of a server failure, you can restore your data onto another server with minimum disruption to users.
- Link aggregation allows you to add performance and fault tolerance to your network connections.

High Availability Hardware Solutions

Apple provides hardware solutions that offer increased reliability, and hence availability:

- Xserve has several features that ensure its availability, including built-in RAID mirroring (allows fast recovery in case of drive failure) and Error Correcting Code (ECC) memory, which helps prevent data corruption.
- Xserve RAID has several built-in redundant components that allow the system to function in case one of them fails. For example, Xserve RAID has two power supplies. If one of the power supplies fails, the system stays on.
- Third-party load-balancing devices allow you to spread the network load throughout a server farm and thus provide scalability and fault tolerance.

This chapter describes the IP failover service and how to configure it in Mac OS X Server.


IP failover is a technology that allows you to set up two computers in a master-backup relationship such that if the master computer fails, the backup computer transparently assumes the role of the master with minimal disruption in service.

For example, if you have a home directory server with 1,000 users and you don't have a backup server, your users will not be able to access their files if the directory server fails. But if you set up another server as a backup, then even if the master server fails the users will be able to access their files through the backup server without being aware of the service disruption, as long as the data is stored on shared storage that is accessible by both computers.

Mac OS X Server provides built-in support for IP failover. In this chapter, you will learn how IP failover works in Mac OS X Server and how to configure it.

IP Failover Overview

IP failover lets you ensure high availability of your servers. A simple IP failover solution consists of two Mac OS X Server computers: a master and a backup. The master computer provides services while the backup computer waits in the background to take over if the master fails.



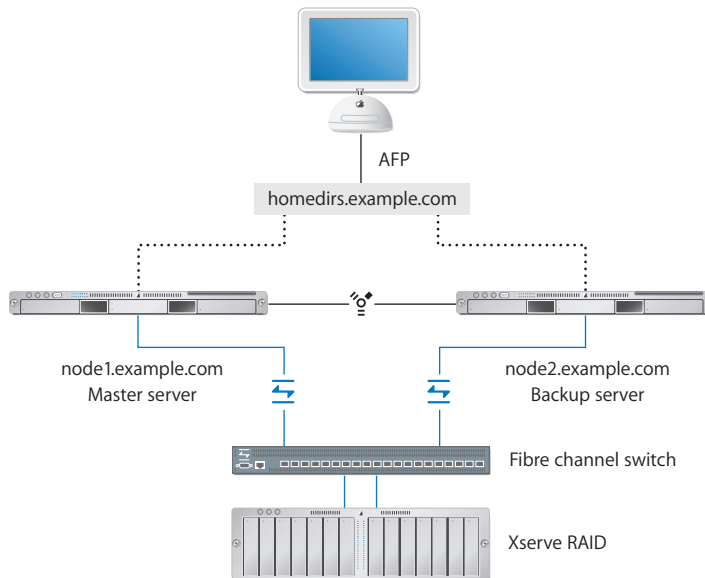
In this scenario, both computers connect to the same network that client requests come over. Each computer has a unique IP address and, optionally, a DNS or domain name. The computers are connected directly to each other using IP over FireWire.

To provide IP failover support, Mac OS X Server uses the `heartbeatd` and `failoverd` daemons:

- The `heartbeatd` daemon runs on the master computer and broadcasts heartbeat messages every second on port 1694 from both network interfaces, announcing the host's availability to other nodes listening with `failoverd`. Sending the heartbeat message on both primary and secondary network links helps prevent false alarms. `heartbeatd` uses the `FAILOVER_BCAST_IPS` entry in the `/etc/hostconfig` file to determine who to send to broadcast the heartbeat message to.
- The `failoverd` daemon runs on the backup computer and listens on port 1694 for broadcasts from a specific address on both interfaces. If `failoverd` stops receiving the heartbeat messages on both interfaces, it takes over the public IP address of the master server, which allows the failover server to service incoming client requests, thus maintaining availability. `failoverd` uses the `FAILOVER_PEER_IP` and `FAILOVER_PEER_IP` entries in the `/etc/hostconfig` file to get the public and private IP addresses of the master server.

For IP failover to work, you must keep the backup computer in sync with the master. For example, if you're using the master computer as an AFP file server, you should make sure that the backup computer has the same AFP service settings. If the settings are not the same, users might not be able to access the file service.

In addition, for IP failover to work, the backup computer should have access to the data that is needed by the client computers. To ensure data availability, you'll have to keep the data on both computers in sync using the `cron` and `rsync` commands or other third-party solutions. Alternatively, you could use shared storage like Xserve RAID in which to store data.



To take advantage of Xserve RAID in IP failover situations, you can use Xsan or third-party storage area network (SAN) software to allow your master and backup servers to access the same volume without corrupting it.

You can also use logical unit number (LUN) masking at the Fibre switch level to grant access to shared data on an Xserve RAID. You'll have to create scripts to instruct the switch to swap access from one server to the other. LUN masking at the switch level ensures that only one server has access to the data, but never both servers at the same time.

Warning: Giving two servers access to the same Xserve RAID volume without Xsan or third party SAN software can corrupt the volume.

After starting the `heartbeatd` and `failoverd` daemons, the master computer starts sending heartbeat messages to the backup server at predetermined intervals. If the backup computer stops receiving these messages, it triggers a chain of events that ends with the backup server taking over the IP address of the master server and assuming its role.

From a client perspective, failover happens transparently, with minimal disruption of service. This is because the client accesses services using a virtual IP address (that is, an address not associated with a particular computer) and/or domain name (for example, `homedirs.example.com`). When the backup server assumes the role of the master, the client doesn't see any difference, as long as services on both computers are configured exactly the same way. A brief disruption of service may be noticeable if IP failover happens while the client is actively communicating with a service. For example, if a user is copying a file from the server and it fails over, the copying process will be disrupted and the user will have to start the copying process again.

Acquiring Master Address—Chain of Events

When the master server fails over, the following chain of events occurs on the backup server:

- 1 The `failoverd` daemon (located in `/usr/sbin/`) detects no broadcasts from the primary server on the FireWire interface.
- 2 The `failoverd` daemon instructs the `NotifyFailover` script (located in `/usr/libexec/`) to send emails to the users listed in `/etc/hostconfig`. If no recipient is specified, an email is sent to the root user.
- 3 `failoverd` executes the `ProcessFailover` script (located in `/usr/libexec/`).
- 4 The `ProcessFailover` script executes the `/Library/Failover/IP_address/Test` script, where `IP_address` is the IP address or domain name of the master server.
 - a If the `Test` script returns false, `ProcessFailover` quits and the backup server does not acquire the IP address of the master server.
 - b If the `Test` script returns true (or does not exist), `ProcessFailover` continues its execution.

Note: By default, the `Test` script is empty, but you can customize it to suit your needs.

- 5 The `ProcessFailover` script executes, in alphabetical order, any script with the prefix `PreAcq` in the `/Library/Failover/IP_address` folder.

`PreAcq` scripts prepares the backup server to acquire the IP address of the master server. By default, Mac OS X Server ships with a number of `PreAcq` scripts, but you can always customize them or add you own.

- 6 The `ProcessFailover` script configures the network interface to use the IP address of the master server.
- 7 The `ProcessFailover` script executes, in alphabetical order, any script with the prefix `PostAcq` in the `/Library/Failover/IP_address` folder.

`PostAcq` scripts run after the backup server acquires the master's IP address. As with `PreAcq` scripts, Mac OS X Server ships with a number of `PostAcq` scripts. But you can add your own. For example, a `PostAcq` script can send you an email to let you know that failover completed successfully.

For more information about `failoverd`, `NotifyFailover`, and `ProcessFailover`, see the corresponding man page or the high availability chapter of the command-line administration guide.

Note: It takes approximately 30 seconds for failover to complete.

Releasing Master Address—Chain of Events

When you trigger failback, the following occurs on the backup server:

- 1 The `failoverd` daemon instructs the `NotifyFailover` script (located in `/usr/libexec/`) to send emails to the users listed in `/etc/hostconfig`. If no recipient is specified, an email is sent to the root user.
- 2 `failoverd` executes the `ProcessFailover` script (located in `/usr/libexec/`).
- 3 The `ProcessFailover` script executes the `Test` script (located in `/Library/Failover/IP_address`, where `IP_address` is the IP address or domain name of the master server).
 - a If the `Test` script returns false, `ProcessFailover` quits and the master server does not acquire the IP address of the backup server.
 - b If the `Test` script returns true (or does not exist), `ProcessFailover` continues its execution.
- 4 The `ProcessFailover` script executes, in alphabetical order, any script with the prefix `PreRel` in the `Library/Failover/IP_address` folder.
- 5 The `ProcessFailover` script releases the IP address of the master server.
- 6 The `ProcessFailover` script executes, in alphabetical order, any script with the prefix `PostRel` in the `Library/Failover/IP_address` folder.

Note: By default, the `Test` script is empty, but you can customize it to suit your needs.

IP Failover Setup

Here is an overview of what you need to do to set up your Mac OS X Server computers for IP failover:

Step 1: Connect the master and backup computers to the same network and configure their TCP/IP settings

This step allows your servers to communicate with the client computers. Each server must have its own unique IP address.

Step 2: Connect the two computers directly using a secondary Ethernet interface or IP over FireWire and configure the IP settings

This step allows direct communication of IP failover events between the servers.

Step 3: Configure and start IP failover service on the master and backup servers

This step allows the master and backup computers to fail over and fail back.

Connecting the Master and Backup Servers to the Same Network

The first step in setting up your servers for failover is to connect the master and backup computers to the same network and configure their network settings.

To connect the master and backup server to the same network:

- 1 Using the primary Ethernet interface, connect the master and backup servers to the same network.
- 2 In the Network pane of System Preferences, configure the TCP/IP settings on both the master and backup computers so that each computer has a unique IP address and both are on the same subnet.

Ideally, you should ask your system administrator to map the IP address of the master server to a virtual DNS name (for example, `homedirs.example.com`) that users use to connect to your server. This allows you to change the IP address of the master server transparently to users.

You might also want to map the IP addresses of the master and backup servers to DNS names (for example, `node1.example.com` and `node2.example.com`) that you can use to refer to the two computers when setting up IP failover.

Connecting the Master and Backup Servers Together

Connect the master and backup computers together using a secondary Ethernet interface or IP over FireWire. This is an important step because the two computers communicate failover events over this connection.

To connect the master and backup servers together:

- 1 Use an Ethernet cable or a FireWire cable to connect the master and backup computers together.
- 2 In the Network pane of System Preferences, configure the TCP/IP settings of the secondary Ethernet interface or IP over FireWire interface on both computers.

Assign each computer a private network IP address (for example `10.1.0.2` and `10.1.0.3`) and make sure that both computers are on the same subnet.

Configuring the Master Server for Failover

Configuring the master server for IP failover is simple. All you have to do is add or edit two entries in the `/etc/hostconfig` file and then restart the server.

To configure the master server for IP failover:

- 1 Add or edit the `FAILOVER_BCAST_IPS` entry in `/etc/hostconfig` to specify the addresses to which to send heartbeat messages.

For example, if the primary IP address of the master server is 171.0.50 and the secondary IP address is 10.1.0.2, you can add the following line to `/etc/hostconfig` to broadcast the message over the two networks:

```
FAILOVER_BCAST_IPS="10.1.0.255 17.1.0.255"
```

However, it's more efficient for your network switch to send the heartbeat messages to specific addresses:

```
FAILOVER_BCAST_IPS="10.1.0.3 17.1.0.51"
```

This line instructs the master server to send the heartbeat messages to the primary and secondary IP addresses of the backup server.

Note: To edit the `/etc/hostconfig` file, you must be root. Use the `sudo` command when opening this file using your preferred command-line editor.

- 2 Add or edit the `FAILOVER_EMAIL_RECIPIENT` entry to specify the email address to send notifications to.

If you don't add this entry, email notifications will go to root.

- 3 Restart the server.

The `IPFailover` startup item launches `heartbeatd` during startup. Upon launch, `heartbeatd` checks its argument list, moves to the background, and periodically sends out "heartbeat" messages to the addresses specified in the `FAILOVER_BCAST_IPS` entry in `/etc/hostconfig`.

Configuring the Backup Server for Failover

Configuring the backup server for IP failover is simple. All you have to do is add or edit two entries in the `/etc/hostconfig` file, disconnect the master and backup servers, restart the backup server, and reconnect the servers.

To configure the backup server for IP failover:

- 1 Add or edit the `FAILOVER_PEER_IP_PAIRS` entry in `/etc/hostconfig` to specify the IP address of the primary network interface on the master server.

For example, if the IP address of the primary network interface on the master server is 171.0.50, add the following entry:

```
FAILOVER_PEER_IP_PAIRS="en0:17.1.0.50"
```

Note: To edit the `/etc/hostconfig` file, you must be root. Use the `sudo` command when opening this file using your preferred command-line editor.

- 2 Add or edit the `FAILOVER_PEER_IP` entry in `/etc/hostconfig` to specify the IP address of the secondary network interface on the master server.

For example, if the IP address of the FireWire port on the master server is 10.1.0.2, add the following entry:

```
FAILOVER_PEER_IP="10.1.0.2"
```

- 3 Disconnect the direct connection between backup server and the master server.

If you're using IP over FireWire for the secondary interface, disconnect the FireWire cable connecting the two computers.

- 4 Restart the backup server.
- 5 Once the backup server has started up, reconnect it to the primary server.

Viewing the IP Failover Log

Mac OS X Server records all IP failover activity in `/Library/Logs/failoverd.log`.

To view failover service log files:

- 1 Open `/Applications/Utilities/Console`.
- 2 Choose `File > Open`.
- 3 Locate and select the `failoverd.log` in the `/Library/Logs/` folder.
- 4 Click `Open`.

You can use the Filter field to display only the log entries you're interested in.

From the Command Line

You can also view the `failoverd.log` file using commands in Terminal. To automate log monitoring, consider using `cron` and `grep` to automatically search log files for IP failover-related keywords and email yourself those entries. For more information, see the IP failover chapter of the command-line administration guide.

Configuring Other Aspects of High Availability

3

This chapter describes how you can improve the availability of your Mac OS X Server solutions.

Eliminating single points of failure and using Xserve and Xserve RAID are some of the things that can boost your server availability. Other things you can do range from simple solutions like using power backup, automatic reboot, and ensuring proper operational conditions (for example, adequate temperature and humidity levels) to more advanced solutions involving link aggregation, load balancing, Open Directory replication, and data backup.

Eliminating Single Points of Failure

To improve the availability of your server, you need to reduce or eliminate single points of failure. A single point of failure is any component in your server environment that, if it fails, causes your server to fail.

Some single points of failure include:

- Computer system
- Hard disk
- Power supply

While it is almost impossible to eliminate all single points of failure, you should at least minimize them as much as possible. For example, using a backup system and the IP failover in Mac OS X Server, as described in Chapter 2, “Configuring IP Failover,” eliminates the computer as a single point of failure. While both the master and backup computers can fail at once or one after the other, the possibility of such an event happening is negligible.

Another way to prevent a computer from failing is to use a backup power source and take advantage of hardware RAID to mirror the hard disk. With hardware RAID, if the main disk fails, the system can still access the same data on the mirror drive, as is the case with Xserve.

Using Xserve and Xserve RAID

Both Xserve and Xserve RAID are designed for extra reliability and hence, high availability.

Using Xserve

While you can use desktop systems like the Power Mac G5 to provide Mac OS X Server services very reliably, the Xserve has several more features that make it ideal for high availability situations.

Some of these features include the following:

- Xserve has eight fans. In the case of a single fan failure, the other fans speed up to compensate, allowing your server to keep running.
- An independent drive architecture isolates the drives electrically, preventing a single drive failure from causing unavailability or performance degradation of the surviving drives—a common problem with multidrive SCSI implementations.
- Xserve G5 uses Error Correction Code (ECC) logic to protect the system from corrupt data and transmission errors. Each DIMM has an extra memory module that stores checksum data for every transaction. The system controller uses this ECC data to identify single-bit errors and corrects them on the fly, preventing unplanned system shutdowns. In the rare event of multiple-bit errors, the system controller detects the error and triggers a system notification to prevent bad data from corrupting further operations. You can set the Server Monitor software to alert you if error rates exceed the defined threshold.
- Xserve G5 has built-in hardware RAID mirroring, which protects your server from failing if the main drive fails.

For more information about Xserve RAID, visit www.apple.com/xserve/.

Using Xserve RAID

Xserve RAID provides you with highly reliable shared storage that is ideal for high availability solutions, especially those involving IP failover.

Xserve RAID reliability/high availability features include these:

- Midplane with a passive data path. The Xserve RAID architecture is designed to avoid vulnerability to a single point of failure. Xserve RAID is built around a midplane with a passive data path, a feature not commonly found in other storage systems of its kind. The midplane is the central connector between the drives, RAID controllers, power supplies, and cooling modules. Most RAID systems depend on the midplane to relay data and instruction sets between drives, and a failure in the midplane can impair data availability. In Xserve RAID, all data passes through the independent drive channels, which are simply held in place by the midplane.

- Dual independent RAID controllers. Two independent storage processor units manage RAID functions, data transfers, and failure protection for each set of seven drives. These two controllers are not redundant, but let you mirror data such that one half of the Xserve RAID can be a mirror for the other. If one controller fails, you can still access data through the second controller.
- Redundant hot-swap cooling modules. The two redundant hot-swap cooling modules provide automatic front-to-rear cooling for rack environments. If one module fails, you can replace it without shutting down the system.
- Power supplies. Either of the two load-sharing and hot-swappable power supplies can power Xserve RAID alone in case one fails.
- Hot sparing. For each RAID controller, any drives not assigned to an array are automatically used as global hot spares. If a drive fails, the RAID controller can automatically rebuild its data on the spare drive without requiring intervention by the administrator. The rebuild operation occurs in the background while the controller processes normal host reads and writes, so that service continues uninterrupted. This gives the administrator ample time to replace the failed drive. Xserve RAID automatically configures the drive as a new hot spare for the array.

For more information about Xserve RAID, visit www.apple.com/xserve/raid/.


Using Backup Power

In the architecture of a server solution, power is a single point of failure. If power goes out, your servers go down without warning. To prevent a sudden disruption in services, you should consider adding a backup source of power. Depending on your application, you might choose to use a standby electrical generator or Uninterruptible Power Supply (UPS) devices to at least gain enough time to notify users of an impending shutdown of services.

Using UPS with Xserve RAID

Xserve RAID provides built-in UPS monitoring via the serial port. In addition to connecting Xserve RAID's power supply to a UPS device, you can connect the two devices together through the serial port, which allows Xserve RAID to monitor the UPS device. When Xserve RAID detects that power is about to go out on the UPS device, it automatically changes the cache mode from high-throughput writeback to safer write-through. This reduces overall system throughput, and thus power consumption, and protects transactions in case UPS power drains out completely.

Note: If you use an APC UPS (a commonly used brand), Xserve RAID notifies you via email when 20% power remains to give you time to gracefully shut down the Xserve RAID.




In addition to UPS, you can also use the optional cache backup batteries with Xserve RAID. These batteries maintain the current transaction data in the cache for up to 72 hours, which gives you ample time to shut down the unit gracefully. When you restore power, Xserve RAID writes the saved transactions in the cache to disk, thus maintaining data integrity, and begins charging the batteries.

While charging the cache backup batteries, Xserve RAID stays in write-through cache mode until the battery charge is over 50%. Then Xserve RAID switches back to write-back mode.

Warning: If not using UPS, you should turn off Xserve RAID's write cache (Xserve RAID does that automatically when connected to UPS) to prevent data loss. The cache backup batteries will back up only the controller cache and not the write cache.

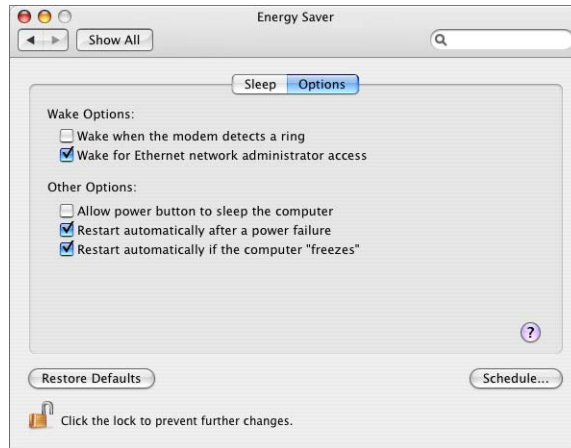
Using UPS With Xserve

Xserve does not provide serial port connectivity to UPS, but can monitor UPS power through the network if the UPS unit has a management network card. Check with UPS vendors for more information.



Setting Up Your Server for Automatic Reboot

You can set up Energy Saver options on your Mac OS X Server computer to automatically restart if it goes down due to a power failure or system freeze.



The automatic reboot options are:

- Restart automatically after a power failure. The power management unit automatically starts up the server after a power failure.
- Restart automatically if the computer “freezes.” The power management unit automatically starts up the server after the server stops responding, has a kernel panic, or freezes.

When you select the option to restart automatically after a “freeze,” Mac OS X Server spawns the `wdticklerd` daemon, which every 30 seconds commands your computer to reboot after 5 minutes. Each time the command is sent, the restart timer is reset. Thus, the timer won’t reach 5 minutes as long as the server is running. But if the computer freezes, the power management unit will restart it after 5 minutes.

To enable automatic reboot:

- 1 Log in to the server as an administrator.
- 2 Open System Preferences and click Energy Saver.
- 3 Click Options.
- 4 Under Other Options, select the two restart options.
- 5 Close System Preferences.

Ensuring Proper Operational Conditions

One of the factors that can cause your servers to malfunction is overheating. This is especially a problem when you cluster computers together in a small space. Other factors such as humidity and power surges can also adversely impact your server.

To protect your servers, you need to make sure that you house them in a place where you can control these factors and provide ideal operating conditions. Check the electrical and environmental requirements for your systems to find what these conditions are.

In addition, make sure that the facility in which you deploy your server has a fire alarm, and prepare a contingency plan to deal with this risk.

Open Directory Replication

If you plan to provide Open Directory services, you should consider creating replicas of your Open Directory master. If the master server fails, client computers can access the replica. For more information, see the section on setting up Open Directory replicas in the Open Directory administration guide.

Protecting Server Security

One of the main threats to the availability of servers is unauthorized access by attackers.

To protect your servers from software attacks, do the following:

- **Enable firewall service.** Your server's firewall is the first line of defense against unauthorized access. For more information, see the chapter on setting up firewall service in the network services administration guide. Consider also a third-party hardware firewall as an additional line of defense if your server is highly prone to attack.
- **Disable unneeded services.** Turning on a service opens up a port from which users can access your system. While enabling firewall service helps fend off unauthorized access, an open port remains a vulnerability that an attacker might be able to exploit.
- **Install antivirus software.** While viruses are far less prevalent on the Mac platform compared with Windows, viruses still pose a risk.
- **Configure Service Access Control Lists (SACLs).** Use SACLs to specify who can access services.
- **Configure Access Control Lists (ACLs).** Use ACLs to control who can access share points and their contents.

To protect your server hardware from theft or sabotage, consider the following:


- House your server in a secure facility. Only authorized personnel should have access to the servers.
- Administer the servers remotely. Manage your servers remotely using applications like Server Admin, Server Monitor, RAID Admin, and Apple Remote Desktop. Minimizing physical access to the systems reduces the possibility of mischief.
- Use security locks. Locking your systems is another prudent thing to do.

Link Aggregation

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

Link aggregation allows you to aggregate or combine multiple physical links connecting your Mac to a link aggregation device (a switch or another Mac) into a single logical link. The result is a fault-tolerant link with a bandwidth equal to the sum of the bandwidths of the physical links.

For example, you can set up an Xserve with four 1-Gbit/s ports (en1, en2, en3, and en4) and use the Network pane of System Preferences to create a link aggregate port configuration (bond0) that combines en1, en2, en3, and en4 into one logical link. The resulting logical link will have a bandwidth of 4 Gbit/s. This link will also provide fault tolerance. If one or more physical links fail, your Xserve's bandwidth will shrink, but the Xserve will still be able to service requests, as long as not all physical links fail at once.



Link aggregation also allows you to take advantage of existing or inexpensive hardware to increase the bandwidth of your server. For example, you can form a link aggregate from a combination of multiple 100-Mbit/s links or 1-Gbit/s links.

The Link Aggregation Control Protocol (LACP)

IEEE 802.3ad Link Aggregation defines a protocol called Link Aggregation Control Protocol (LACP) and is used by Mac OS X Server to aggregate (combine) multiple ports into a link aggregate (a virtual port) that can be used for TCP and UDP connections. When you define a link aggregate, the nodes on each side of the aggregate (for example, a computer and a switch) use the LACP protocol over each physical link to:

- Determine whether the link can be aggregated.
- Maintain and monitor the aggregation.

If a node doesn't receive LACP packets from its peer (the other node in the aggregate) regularly, it assumes that the peer is no longer active, and automatically removes the port from the aggregate.

In addition to LACP, Mac OS X Server uses a frame distribution algorithm to map a "conversation" to a particular port. This algorithm sends packets to the system on the other end of the aggregate only if it has packet reception enabled. In other words, the algorithm won't send packets if the other system isn't listening. Mapping a "conversation" to a particular port guarantees that packet reordering will not occur.

Link Aggregation Scenarios


Following are three common aggregation scenarios that you can set up:

- Computer-to-computer
- Computer-to-switch
- Computer-to-switch-pair

These scenarios are described in the following sections.


Computer-to-Computer

In this scenario, you connect the two servers directly using the physical links of the link aggregate. This allows the two servers to communicate at a higher speed without the need for a switch. This configuration is ideal for ensuring back-end redundancy.



Computer-to-Switch

In this scenario, you connect your server to a switch configured for 802.3ad link aggregation. The switch should have a bandwidth for handling incoming traffic equal to or greater than that of the link aggregate (logical link) you define on your server.




For example, if you create an aggregate of four 1-Gbit/s links, you should use a switch that can handle incoming traffic (from clients) at 4 Gbit/s or more. Otherwise, the increased bandwidth advantage in the link aggregate won't be fully realized.

Note: For information on how to configure your switch for 802.3ad link aggregation, refer to the documentation provided by the switch manufacturer.

Computer-to-Switch-Pair

In this scenario, you improve on the computer-to-switch scenario by using two switches to eliminate the switch as a single point to failure.




For example, you can connect two links of the link aggregate to the master switch and the remaining links to the backup switch. As long as the master switch is active, the backup switch remains inactive. If the master switch fails, the backup switch takes over automatically and transparently to the user.

While this scenario adds redundancy that protects the server from becoming unavailable if the switch fails, it results in decreased bandwidth.

Setting Up Link Aggregation in Mac OS X Server

To set up your Mac OS X Server for link aggregation, you need a Mac with two or more IEEE 802.3ad-compliant Ethernet ports. In addition, you need at least one IEEE 802.3ad-compliant switch or another Mac OS X Server computer with the same number of ports.

You create a link aggregate on your computer in the Network pane of System Preferences.



To create a link aggregate:

- 1 Log in to the server as an administrative user.
- 2 Open System Preferences.
- 3 Click Network.
- 4 Choose Network Port Configurations from the Show pop-up menu.
- 5 Click New.
- 6 Choose Link Aggregate from the Port pop-up menu.
Note: You'll only see this option if you have two or more Ethernet interfaces on your system.
- 7 Type the name of the link aggregate in the Name field.
- 8 Select the ports to aggregate from the list.
- 9 Click OK.
- 10 Click Apply Now.

By default the system gives the link aggregate the interface name `bond<num>`, where `<num>` is a number indicating precedence. For example, the first link aggregate will be named `bond0`, the second `bond1`, and the third `bond2`.

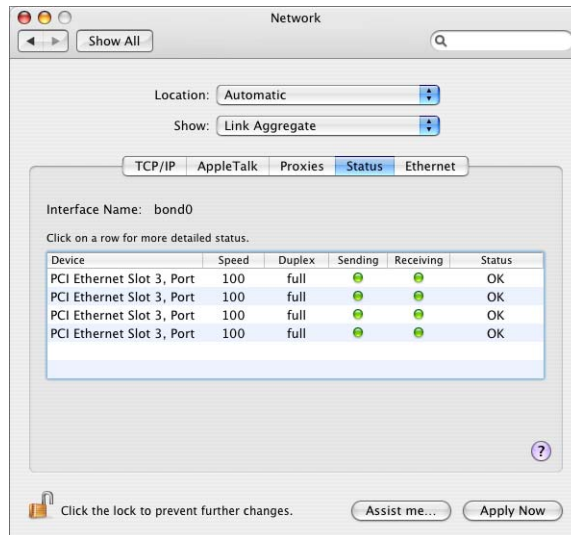
The interface name `bond<num>` assigned by the system is different from the name you give to the link aggregate port configuration. The interface name is for use in the command line, while the port configuration name is for use in the Network pane of System Preferences.

For example, if you enter the command `ifconfig -a`, the output refers to the link aggregate using the interface name and *not* the port configuration name:

```
...
bond0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    inet6 fe80::2e0:edff:fe08:3ea6 prefixlen 64 scopeid 0xc
    inet 10.0.0.12 netmask 0xffffffff broadcast 10.0.0.255
    ether 00:e0:ed:08:3e:a6
    media: autoselect (100baseTX <full-duplex>) status: active
    supported media: autoselect
    bond interfaces: en1 en2 en3 en4
```

Monitoring Link Aggregation Status

You can monitor the status of a link aggregate in Mac OS X and Mac OS X Server using the Status pane of the Network pane of System Preferences.



To monitor the status of a link aggregate:

- 1 Open System Preferences.
- 2 Click Network.
- 3 Choose the link aggregate port configuration from the Show pop-up menu.
- 4 Click Status.

The Status pane displays a list containing a row for each physical link in the link aggregate. For each link, you can view the name of the network interface, its speed, its duplex setting, the status indicators for incoming and outgoing traffic, and an overall assessment of the status.

Note: The Sending and Receiving status indicators are color-coded. Green means the link is active (turned on) and connected. Yellow means the link is active but not connected. Red means the link can't send or receive traffic.


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

Load Balancing

One of the factors that can cause your services to become unavailable is server overload. A server has limited resources and can service a limited number of requests simultaneously. If the server gets overloaded, it will slow down and might eventually crash.

One way to overcome this problem is to distribute the load among a group of servers (a server farm) using a third-party load-balancing device. Clients send requests to the device, which then forwards the request to the first available server based on a pre-defined algorithm. The clients see only a single virtual address, that of the load-balancing device.

Many load-balancing devices also function as switches, providing two functions in one, which reduces the amount of hardware you need to use.



Note: A load-balancing device must be able to handle the aggregate (combined) traffic of the servers connected to it. Otherwise, the device itself becomes a bottleneck, which reduces the availability of your servers.

Load balancing provides several advantages:

- **High availability.** Distributing the load among multiple servers helps you reduce the chances that an individual server will fail due to server overload.
- **Fault tolerance.** If a server fails, traffic is transparently redirected to other servers. There might be a brief disruption of service if, for example, a server fails while a user is downloading a file from shared storage, but the user will be able to reconnect and restart the file download process.
- **Scalability.** If demand for your services increases, you can transparently add more servers to your farm to keep up with the demand.
- **Better performance.** By sending requests to the least-busy servers, you can respond faster to user requests.

Backup

Sometimes, a server failure or loss of data is inevitable. But what you can do is back up important data so you can restore your server to operating condition with a minimal down time if necessary.

Developing a Backup Strategy

A good backup strategy is comprehensive and takes into account all the aspects that guarantee successful data recovery. A comprehensive backup strategy includes:

- Defining backup scope
- Choosing a backup method
- Choosing a backup rotation scheme
- Choosing media type
- Defining a backup verification mechanism

Defining Backup Scope

Before you determine the best way to back up and restore data, you need to define what you need to back up. For example, you might want to create an image of an entire system. When the main system fails, you can restore the system from the image instead of having to restore specific files scattered around the system. There are cases, however, in which you're only interested in backing up certain files, such as home directories and configuration files.

Whether you want to back up an entire system or a small set of files, defining the backup scope helps you determine how much backup storage you need and how often. Data that changes frequently needs to be backed up more often than data that changes only once or twice a month.

Choosing a Backup Method

A backup method defines what you back up. There are three common backup methods:

- **Full backup.** In a full backup, you copy all the files, even if they haven't changed since the last backup. The backup and restoration procedures are simple, but the storage requirement is high.
- **Differential backup.** You start with a full backup. Then, you back up only the files that have changed since the last "full" backup. This means that a file that changed only once after the full backup will be backed up every time you perform a differential backup, even if it had not changed since. The differential backup and restoration procedures are complex because they need to determine which files changed, but the storage requirements are medium.
- **Incremental backup.** You start with a full backup. Then, you back up only the files that changed since the last backup. This method is the most storage-efficient method, but like differential backup, it is more complex than the full backup method.

Choosing a Backup Rotation Scheme

A backup rotation scheme determines the most efficient way to back up data over a specific period of time. An example of a rotation scheme is the grandfather-father-son rotation scheme. In this scheme, you perform incremental daily backups (son), and full weekly (father) and monthly (grandfather) backups.

In the grandfather-father-son rotation scheme, the number of media sets you use for backup determines how much backup history you have. For example, if you use eight backup sets for daily backups, you have eight days of daily backup history because you'll recycle media sets every eight days.

Choosing Media Type

There are several factors that help you determine what type of media to choose:

- **Cost.** Use the cost per GB to determine what media to choose. For example, if your storage needs are limited, you can justify higher cost per GB. But if you need a large amount of storage, cost becomes a big factor in your decision. One of the most cost-effective storage solutions is Xserve RAID. Not only does it provide you with a low cost per GB, but also it doesn't require the special handling needed by other cost-effective storage types, such as tape drives.
- **Capacity.** If you back up only a small amount of data, low-capacity storage media can do the job. But if you need to back up large amounts of data, use high-capacity devices, such as Xserve RAID.
- **Speed.** When your goal is to keep your server available most of the time, restoration speed becomes a big factor in deciding which type of media to choose. Tape backup systems can be very cost-effective, but are much slower than Xserve RAID.
- **Reliability.** Successful restoration is the goal of a good backup strategy. If you can't restore lost data, all the effort and cost you spent in backing up data is wasted and the availability of your services compromised. Therefore, it's important that you choose highly reliable media to prevent data loss. For example, tapes are more reliable than hard disks because they don't contain moving parts.
- **Archive life.** You never know when you'll need your backed up data. Therefore, you must choose media that is designed to last for a long time. Dust, humidity, and other factors can damage storage media and result in data loss.

Defining a Backup Verification Mechanism

A backup is no good if you can't use it to restore lost data. You should have a strategy for regularly conducting test restorations. Some third-party software providers support this functionality. But if you're using your own backup solutions, you need to develop the necessary test procedures.

Command-Line Backup and Restoration Tools

Mac OS X Server provides several command-line tools for data backup and restoration:

- `rsync`. Use this command to keep a backup copy of your data in sync with the original. `rsync` only copies the files that have changed and doesn't copy resource forks.
- `ditto`. Use this command to perform full backups.
- `asr`. Use this command to back up and restore an entire volume.

For more information about these commands, refer to the command-line administration guide.

Note: You can use the `cron` command to automate data backup using the aforementioned commands. For more information about using `cron`, refer to the command-line administration guide.

Effective monitoring allows you to detect potential problems before they occur and gives you early warning when they occur.

Detecting potential problems allows you to take steps to resolve them before they impact availability of your servers. In addition, getting early warning when a problem does occur allows you to take corrective action quickly and minimize disruption to your services.

This chapter briefly describes monitoring tools and tells you where you can find more information:

- Console
- Server Monitor
- RAID Admin
- Disk monitoring tools
- Network monitoring tools

Console

Use Console to monitor relevant log files for potential problems that may cause your server to fail.

For example, you can monitor your web server's `/var/log/httpd/access_log` file for signs of denial of service attacks. If you detect these signs, you can immediately implement a preplanned response to prevent your web server from becoming unavailable.

To improve your log monitoring efficiency, consider automating the monitoring process using AppleScript or Terminal commands like `grep` and `cron`. Refer to the command-line administration guide for more information about using `grep` and `cron`.

Server Monitor

The Server Monitor application can issue alerts via email, cell phone, or pager notification as soon as it detects critical problems. Built-in sensors detect and report essential operating factors like power, temperature, and condition of several key components.

The Server Monitor user interface gives you the ability to quickly detect problems. In the main window, Server Monitor lists each server on a separate line, along with temperature information and the status of each of its components, including fans, disk drives, memory modules, power supplies, and Ethernet connections. A green status indicator denotes the component is OK, a yellow status indicator denotes a warning, and a red status indicator denotes an error.

For more information about Server Monitor, choose Server Monitor Help from Server Monitor's Help menu.

RAID Admin

Like Server Monitor, you can configure RAID Admin to send you an email or page when a component is in trouble. For every unit, RAID Admin displays the status of the unit and each of its components, including disk drives, fibre channel, and network connections. RAID Admin uses green, yellow, or red status indicators. You can also configure it to send you an email or page when a component is in trouble.

For more information about RAID Admin, choose RAID Admin Help from RAID Admin's Help menu.

In addition, RAID Admin provides you with a bird's eye view of status of the Xserve RAID units that appear in the main window.

Disk Monitoring Tools

Running out of disk space can cause your server to become unreliable and probably fail. To prevent this from happening, you need to constantly monitor disk space usage on your servers and delete or back up files to clear disk space.

Mac OS X Server ships with a number of command-line tools that you can use to monitor disk space on your computer:

- `df`. This command tells you how much space is used and how much is available on every mounted volume.

For example, the following command lists local volumes and displays disk usage in human-readable form:

```
df -Hl
Filesystem      Size  Used Avail Capacity  Mounted on
/dev/disk0s9    40G   38G   2.1G    95%      /
```

In this example, the hard disk is almost full with only 2.1 GB left. This tells you that you should act immediately to free up space on your hard drive before it fills up and causes problems for your users.

- `du`. This command tells you how much space is used by specific folders or files. For example, the following command tells you how much space is used by each user's home folder:

```
sudo du -sh /Users/*
3.2M    /Users/Shared
9.3M    /Users/omar
8.8M    /Users/jay
1.6M    /Users/lili
...
```

Knowing who's consuming most of the space on the hard drive lets you contact users by email and ask them to delete unused files.

Note: With Workgroup Manager, you can set disk quotas for users and generate disk usage reports. Refer to the user management guide for more information.

- `diskspacesmonitor`. This command lets you automate the process of monitoring disk space usage. When the amount of free disk space drops below the level you specify, `diskspacesmonitor` executes shell scripts which send you a notification. This command defines two action levels:
 - Alert—Sends you a warning message when disk space usage reaches 75%.
 - Recover—Archives rarely used files and deletes unneeded files when disk space usage reaches 85%.

For more information about these commands, see the corresponding man page or the high availability chapter of the command-line administration guide.

Network Monitoring Tools

Degradation in network performance or other network problems can adversely affect the availability of your services. Several network monitoring tools can alert you to possible problems early on, so you can take corrective action to avoid or minimize down time.

- To monitor network activity, you can use the `tcpdump` utility in Mac OS X Server. This utility prints out the headers of incoming and outgoing packets on a network interface that match the specified parameters.
- Using `tcpdump` to monitor network traffic is especially useful when trying to detect denial of service attacks. For example, the following command monitors all incoming traffic on port 80 on your computer:

```
sudo tcpdump -i en0 dst port 80
```

If you detect an unusual number of requests coming from the same source, you can use the firewall service to block all traffic from that source.

For more information about `tcpdump`, see the corresponding man page or the high availability chapter of the command-line administration guide.

- Using `tcpdump` to monitor traffic can be time consuming. This is why you should use AppleScripts or shell scripts to automate the monitoring process.
- Also consider using Ethereal, an X11 open source packet sniffing tool that you can run in the X11 environment on Mac OS X Server. This tool, unlike `tcpdump`, has a graphical user interface and a set of powerful network analysis tools. For more information about Ethereal, visit www.ethereal.com/.
- In addition, you can use other third-party tools that automatically analyze network traffic and alert you to problems.

Try these suggestions to solve or avoid failover problems while configuring or using failover service.

Service Doesn't Automatically fail over When Primary Service Is Unavailable

- Verify that the server's software serial number is entered correctly and has not expired. To check the number, open Server Admin, select the server in the Computers & Services list, and click Overview. To enter an updated serial number, click Settings.
- Check the IP failover log (/Library/Logs/failoverd.log) for problem indications.
- Make sure cables are attached correctly to the hardware components.
- Verify that the network settings are configured correctly on both the master and backup server.
- Verify that `FAILOVER_BCAST_IPS` entry in the `/etc/hostconfig` file is correctly configured on the master server.
- Verify that the `FAILOVER_PEER_IP_PAIRS` and `FAILOVER_PEER_IP` entries in the `/etc/hostconfig` file are correctly configured on the backup server.

Email Recipient Isn't Getting Notifications

- Verify that the email addresses specified by the `FAILOVER_EMAIL_RECIPIENT` entry in the `/etc/hostconfig` file is correct.
- Make sure your public network interface is working.
- Make sure you have correctly configured IP failover.

Failover Took Place, But Still Having Problems

- Make sure you shutdown the master server after failover.

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

administrator A user with server or directory domain administration privileges. Administrators are always members of the predefined “admin” group.

AFP Apple Filing Protocol. A client/server protocol used by Apple file service on Macintosh-compatible computers to share files and network services. AFP uses TCP/IP and other protocols to communicate between computers on a network.

aggregation Combining similar objects or resources (such as disks or network connections) into a single logical resource in order to achieve increased performance. For example, two or more disks can be aggregated into a single logical disk to provide a single volume with increased capacity.

Apple Filing Protocol See **AFP**.

AppleScript A scripting language with English-like syntax, used to write script files that can control your computer. AppleScript is part of the Mac operating system and is included on every Macintosh.

automatic backup A backup triggered by an event (such as a scheduled time, or the exceeding of a storage limit) rather than by a human action.

automatic failover Failover that occurs without human intervention.

availability The amount of time that a system is available during those time periods when it’s expected to be available. See also **high availability**.

back up (verb) The act of creating a backup.

backup (noun) A collection of data that’s stored for purposes of recovery in case the original copy of data is lost or becomes inaccessible.

bandwidth The capacity of a network connection, measured in bits or bytes per second, for carrying data.

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

bit rate The speed at which bits are transmitted over a network, usually expressed in bits per second.

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

client A computer (or a user of the computer) that requests data or services from another computer, or server.

cluster A collection of computers interconnected in order to improve reliability, availability, and performance. Clustered computers often run special software to coordinate the computers' activities. See also **computational cluster**.

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

Common Internet File System See **SMB/CIFS**.

computational cluster A group of computers or servers that are grouped together to share the processing of a task at a high level of performance. A computational cluster can perform larger tasks than a single computer would be able to complete, and such a grouping of computers (or "nodes") can achieve high performance comparable to a supercomputer.

daemon A program that runs in the background and provides important system services, such as processing incoming email or handling requests from the network.

data rate The amount of information transmitted per second.

default The automatic action performed by a program unless the user chooses otherwise.

denial of service See **DoS attack**.

denial of service attack See **DoS attack**.

deploy To place configured computer systems into a specific environment or make them available for use in that environment.

disk A rewritable data storage device. See also **disk drive**, **logical disk**.

disk image A file that, when opened, creates an icon on a Mac OS desktop that looks and acts like an actual disk or volume. Using NetBoot, client computers can start up over the network from a server-based disk image that contains system software. Disk image files have a filename extension of either .img or .dmg. The two image formats are similar and are represented with the same icon in the Finder. The .dmg format cannot be used on computers running Mac OS 9.

DNS Domain Name System. A distributed database that maps IP addresses to domain names. A DNS server, also known as a name server, keeps a list of names and the IP addresses associated with each name.

DNS domain A unique name of a computer used in the Domain Name System to translate IP addresses and names. Also called a **domain name**.

DNS name A unique name of a computer used in the Domain Name System to translate IP addresses and names. Also called a **domain name**.

domain Part of the domain name of a computer on the Internet. It does not include the Top Level Domain designator (for example, .com, .net, .us, .uk). Domain name "www.example.com" consists of the subdomain or host name "www," the domain "example," and the top level domain "com."

domain name See **DNS name**.

Domain Name System See **DNS**.

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

Ethernet A common local area networking technology in which data is transmitted in units called packets using protocols such as TCP/IP.

Ethernet adapter An adapter that connects a device to an Ethernet network. Usually called an Ethernet card or Ethernet NIC. See also **NIC**.

fault tolerance The ability of a system to continue to perform its function when one or more of its components has failed.

Fibre Channel The architecture on which most SAN implementations are built. Fibre Channel is a technology standard that allows data to be transferred from one network node to another at very high speeds.

file system A scheme for storing data on storage devices that allows applications to read and write files without having to deal with lower-level details.

GB Gigabyte. 1,073,741,824 (2³⁰) bytes.

Gigabit Ethernet A group of Ethernet standards in which data is transmitted at 1 gigabit per second (Gbit/s). Abbreviated GBE.

gigabyte See GB.

high availability The ability of a system to perform its function continuously (without interruption).

home directory A folder for a user's personal use. Mac OS X also uses the home directory, for example, to store system preferences and managed user settings for Mac OS X users.

host name A unique name for a server, historically referred to as the UNIX hostname. The Mac OS X Server host name is used primarily for client access to NFS home directories. A server determines its host name by using the first name available from the following sources: the name specified in the /etc/hostconfig file (HOSTNAME=some-host-name); the name provided by the DHCP or BootP server for the primary IP address; the first name returned by a reverse DNS (address-to-name) query for the primary IP address; the local hostname; the name "localhost."

hot spare A spare disk that's running and ready to be written to, and that a RAID system can use instantly in place of a failed disk.

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

Hypertext Transfer Protocol See HTTP.

image See disk image.

Internet Generally speaking, a set of interconnected computer networks communicating through a common protocol (TCP/IP). The Internet (note the capitalization) is the most extensive publicly accessible system of interconnected computer networks in the world.

Internet Protocol See IP.

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.

KB Kilobyte. 1,024 (2¹⁰) bytes.

kilobyte See KB.

link An active physical connection (electrical or optical) between two nodes on a network.

link aggregation Configuring several physical network links as a single logical link to improve the capacity and availability of network connections. With link aggregation, all ports are assigned the same ID. Compare to **multipathing**, in which each port keeps its own address.

load balancing The process of distributing client computers' requests for network services across multiple servers to optimize performance.

log in (verb) The act of starting a session with a system (often by authenticating as a user with an account on the system) in order to obtain services or access files. Note that logging in is separate from connecting, which merely entails establishing a physical link with the system.

logical unit number See **LUN**.

LUN Logical unit number. A SCSI identifier for a logical storage device. In Xsan, an unformatted logical storage device such as an Xserve RAID array or slice.

LUN mapping A way to avoid simultaneous writes to an array by assigning a specific LUN to a single host computer. LUN mapping happens at the RAID level. Compare to LUN masking.

LUN masking A way to avoid simultaneous writes to an array by "hiding," or masking access to, a LUN for unwanted host computers within a specific zone. LUN masking is done at the Fibre Channel switch. Compare to LUN mapping.

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.

MB Megabyte. 1,048,576 (2^{20}) bytes.

MB/s Abbreviation for megabytes per second.

Mbit Abbreviation for megabit.

Mbit/s Abbreviation for megabits per second.

media The material in a storage device on which data is recorded.

megabyte See **MB**.

mirrored Refers to a disk array that uses RAID 1, or mirroring.

mirroring Writing identical copies of data to two physical drives. Mirroring protects data against loss due to disk failure, and is the simplest method of achieving data redundancy.

name server A server on a network that keeps a list of names and the IP addresses associated with each name. See also **DNS**, **WINS**.

Network File System See **NFS**.

network interface Your computer's hardware connection to a network. This includes (but isn't limited to) Ethernet connections, AirPort cards, and FireWire connections.

network interface card See **NIC**.

NFS Network File System. A client/server protocol that uses Internet Protocol (IP) to allow remote users to access files as though they were local. NFS exports shared volumes to computers according to IP address, rather than user name and password.

offline Refers to data that isn't immediately available, or to devices that are physically connected but not available for use.

online Refers to data, devices, or network connections that are available for immediate use.

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 source A term for the cooperative development of software by the Internet community. The basic principle is to involve as many people as possible in writing and debugging code by publishing the source code and encouraging the formation of a large community of developers who will submit modifications and enhancements.

port A sort of virtual mail slot. A server uses port numbers to determine which application should receive data packets. Firewalls use port numbers to determine whether data packets are allowed to traverse a local network. "Port" usually refers to either a TCP or UDP port.

port name A unique identifier assigned to a Fibre Channel port.

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

RAID Redundant Array of Independent (or Inexpensive) Disks. A grouping of multiple physical hard disks into a disk array, which either provides high-speed access to stored data, mirrors the data so that it can be rebuilt in case of disk failure, or both of these features. The RAID array is presented to the storage system as a single logical storage unit. See also **RAID array**, **RAID level**.

RAID 1 A RAID scheme that creates a pair of mirrored drives with identical copies of the same data. It provides a high level of data availability.

RAID array A group of physical disks organized and protected by a RAID scheme and presented by RAID hardware or software as a single logical disk. In Xsan, RAID arrays appear as LUNs, which are combined to form storage pools.

RAID level A storage allocation scheme used for storing data on a RAID array. Specified by a number, as in RAID 3 or RAID 0+1.

redundancy The duplication of data, or the inclusion of extra components in a system (such as additional disk drives), in order to recover data or continue the operation of the system after the failure of a system component.

router A computer networking device that forwards data packets toward their destinations. A router is a special form of gateway which links related network segments. In the small office or home, the term router often means an Internet gateway, often with Network Address Translation (NAT) functions. Although generally correct, the term router more properly refers to a network device with dedicated routing hardware.

server A computer that provides services (such as file service, mail service, or web service) to other computers or network devices.

Server Message Block/Common Internet File System See **SMB/CIFS**.

slice A logical subdivision of a RAID array. Each slice is a separate LUN and appears as a separate volume on a host computer.

SMB/CIFS Server Message Block/Common Internet File System. A protocol that allows client computers to access files and network services. It can be used over TCP/IP, the Internet, and other network protocols. Windows services use SMB/CIFS to provide access to servers, printers, and other network resources.

switch Networking hardware that connects multiple nodes (or computers) together. Switches are used in both Ethernet and Fibre Channel networking to provide fast connections between devices.

Index

A

Access Control Lists
See also ACLs 24
 ACLs 24
 antivirus 24
 Apple Remote Desktop 25
 asr 35
 automatic reboot 23

B

backup 33
 asr 35
 cron 35
 ditto 35
 media 34
 rotation schemes 34
 rsync 10, 35
 verification mechanism 34
 backup methods 33

C

Console 37
 cron 35

D

daemons
 failoverd 12
 heartbeatd 12
 launchd 5, 10
 watchdog 5
 differential backup 33
 disk monitoring tools
 df 39
 diskspacemonitor 39
 du 39
 ditto 35
 documentation 6

E

Energy Saver pane 10
 Etherreal 40

F

FAILOVER_BCAST_IPS 12
 FAILOVER_PEER_IP 12
 failoverd 12
 failoverd.log 18
 firewall 24
 FireWire 16
 frame distribution algorithm 27
 full backup 33

G

grep 37

H

heartbeatd 12
 high availability, defined 9
 hostconfig entries
 FAILOVER_BCAST_IPS 12
 FAILOVER_PEER_IP 12

I

incremental backup 33
 IP failover
 backup server 12
 defined 11
 log 18
 master server 12
 setup 15
 troubleshooting 41
 Xserve RAID 13
 IP over FireWire 16

L

LACP 26
 launchd 5, 10
 link aggregation 5, 10
 defined 25
 frame distribution algorithm 27
 monitoring 31
 setup 29
 Link Aggregation Control Protocol
 See also LACP 26

load balancing 32
load-balancing devices 10

M

mirroring 10
monitoring tools
 Apple Remote Desktop 25
 Console 37
 disk monitoring tools 39
 network monitoring tools 40
 RAID Admin 10, 25, 38
 Server Monitor 25, 38

N

network monitoring tools 39
 Ethereal 40
 tcpdump 40
NotifyFailover 14, 15

O

Open Directory 24

P

PostAcq 14
PreAcq 14
ProcessFailover 14

R

RAID 19
RAID Admin 10, 25, 38
reboot, automatic 23
rotation scheme
 grandfather-father-son 34
rsync 10, 35

S

SACLS 24
scripts

NotifyFailover 14, 15
PostAcq 14
PreAcq 14
ProcessFailover 14
Test 14
Security 24
Server Admin 25
server administration guides 6
Server Monitor 10, 25, 38
Service Access Control Lists
 See also SACLS 24
single points of failure 19
switch 28

T

TCP 26
tcpdump 40
Test 14

U

UDP 26
Uninterruptible Power Supply
 See also UPS 21
UPS 21
 APC 22

V

virtual address 32
virtual DNS name 16
virtual port 26

W

watchdog 5

X

Xserve 10, 20
Xserve RAID 10, 13, 20