




Mac OS X Server

Xgrid Administration and
High Performance Computing
For Version 10.5 Leopard

 Apple Inc.
© 2007 Apple 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 Inc. is not responsible for printing or clerical errors.

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

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

AirPort, Apple, the Apple logo, Bonjour, FireWire, iPod, Mac, Macintosh, Mac OS, Xgrid, Xsan, and Xserve are trademarks of Apple Inc., registered in the U.S. and other countries. Apple Remote Desktop and Finder are trademarks of Apple Inc.

Intel, Intel Core, and Xeon are trademarks of Intel Corp. in the U.S. and other countries.

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

UNIX is a registered trademark of The Open Group.

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.

019-0946/2007-09-01

Contents

Preface	9 About This Guide
	9 What's New in Xgrid Administration
	9 What's in This Guide
	10 Using This Guide
	10 Using Onscreen Help
	11 Advanced Server Administration Guides
	12 Viewing PDF Guides on Screen
	12 Printing PDF Guides
	13 Getting Documentation Updates
	13 Getting Additional Information

Part I Xgrid Administration

Chapter 1	17 Introducing Xgrid Service
	17 About Xgrid and Computational Grids
	18 How Xgrid Works
	20 Common Types of Grids and Grid Computing Styles
	20 Xgrid Clusters
	21 Local Grids
	21 Distributed Grids
	22 Xgrid Components
	23 Agent
	23 Client
	24 Controller
	24 Jobs
	24 Requirements and Capacities
Chapter 2	25 Setting Up and Configuring Xgrid Service
	25 Setup Overview
	26 Before Setting Up Xgrid Service
	26 Authentication Methods for Xgrid
	27 Single Sign-On (SSO)

27	Password-Based Authentication
27	No Authentication
28	Hosting the Grid Controller
28	Turning Xgrid Service On
28	Configuring Xgrid with the Xgrid Service Configuration Assistant
29	Configuring Xgrid to Host a Grid Using the Xgrid Service Configuration Assistant
29	Configuring Xgrid to Join a Grid Using Xgrid Service Configuration Assistant
30	Setting Up Xgrid Service
30	Xgrid and Multiple Network Interfaces
30	Configuring Controller Settings
31	Starting Xgrid Service
32	Configuring an Xgrid Agent (Mac OS X Server)
33	Configuring an Xgrid Agent (Mac OS X)
34	Setting Up Grid Authentication
34	Setting Up Kerberos for Xgrid
34	Setting Passwords for Xgrid
35	Managing Client Access
35	Setting SAACL Permissions for Users and Groups
36	Setting SAACL Permissions for Administrators
37	Managing Xgrid Service
37	Viewing Xgrid Service Status
37	Viewing Xgrid Service Logs
38	Stopping Xgrid Service

Chapter 3

39	Managing a Grid
39	Using Xgrid Admin
40	Status Indicators in Xgrid Admin
40	Managing the Xgrid Controller
40	Connecting to an Xgrid Controller
41	Disconnecting from an Xgrid Controller
41	Adding an Xgrid Controller
41	Removing an Xgrid Controller
42	Managing Agents
43	Viewing a List of Agents
43	Adding an Agent
44	Deleting an Agent
44	Managing Jobs
44	Viewing a List of Jobs
44	Stopping a Job
45	Repeating or Restarting a Job
45	Deleting a Job
45	Adding a Grid
46	Deleting a Grid

- 46 Monitoring Grid Activity
- Chapter 4**
- 47 **Planning and Submitting Xgrid Jobs**
- 47 Structuring Jobs for Xgrid
- 47 About Job Styles
- 48 About Job Failure
- 48 Submitting a Job
- 48 Examples of Xgrid Job Submission and Results Retrieval
- 49 Viewing Job Status
- 49 Retrieving Job Results

- Chapter 5**
- 51 **Solving Xgrid Problems**
- 51 If Your Agents Can't Connect to the Xgrid Controller
- 51 If You Use Xgrid over SSH
- 52 If You Run Tasks on Multi-CPU Machines
- 52 If You Submit a Large Number of Jobs
- 53 If You Want to Use Xgrid on Other Platforms
- 53 If the Xgrid Controller Must Be Restarted
- 53 If Xgrid Has Crashed
- 53 If You Are Trying to Submit Jobs over 2 GB
- 54 If You Want to Enable Kerberos/SSO for Xgrid
- 55 For More Information

Part II **Configuring High Performance Computing**

- Chapter 6**
- 59 **Introducing High Performance Computing**
- 59 Understanding HPC
- 59 Apple and HPC
- 60 Mac OS X Server
- 60 Xserve Clusters
- 60 Xserve 64-Bit Architecture
- 62 Support of Loosely Coupled Computations

- Chapter 7**
- 63 **Reviewing the Cluster Setup Process**
- 64 Cluster Setup Overview

- Chapter 8**
- 67 **Identifying Prerequisites and System Requirements**
- 67 Prerequisites
- 67 Expertise
- 67 Xserve Configuration
- 68 System Requirements
- 68 Infrastructure Requirements
- 72 Software Requirements

	72	Private Network Requirements
	73	Static IP Address and Hostname Requirements
Chapter 9	75	Preparing the Cluster for Configuration
	75	Preparing the Cluster Nodes for Software Configuration
	78	(Optional) Setting Up the Management Computer
Chapter 10	81	Setting Up the Cluster Controller
	81	Setting Up Server Software on the Cluster Controller
	84	Configuring DNS Service
	85	Verifying DNS Settings
	86	Configuring Open Directory Service
	86	Configuring the Cluster Controller as an Open Directory Master
	87	Configuring DHCP Service
	88	Configuring Firewall Settings on the Cluster Controller
	90	Configuring NAT Settings on the Cluster Controller
	90	Configuring NFS
	90	Configuring VPN Service
	91	Configuring Xgrid Service
	92	Preparing the Data Drive as a Mirrored RAID set
	93	Creating a Home Directory Automount Share Point
	94	Creating User Accounts
Chapter 11	95	Setting Up Compute Nodes
	95	Creating an Auto Server Setup Record for Compute Nodes
	98	Verifying LDAP Record Creation
	98	Setting Up Compute Nodes
	99	Configuring Cluster Nodes
	101	Creating and Verifying a VPN Connection
	101	Joining a Remote Client to the Kerberos Realm
	102	Verifying Remote Client Access to the Kerberos Realm
Chapter 12	103	Testing Your Cluster
	103	Checking Your Cluster Using Xgrid Admin
	104	Testing Your Xgrid Cluster
	105	Verifying Your Xgrid Configuration
	106	Verifying Your SSH Connection
Appendix A	107	Cluster Setup Checklist
Appendix B	111	Automating Compute Node Configuration
	111	Naming Multiple Cluster Nodes
	112	Joining Multiple Cluster Nodes to the Kerberos Realm
	112	Configuring Xgrid Agent Settings Using Apple Remote Desktop

	114	Using SSH Without Passwords
Glossary	115	
Index	121	

About This Guide

This guide describes the Xgrid components included in Mac OS X Server and tells you how to configure and use them in computational grids.

Xgrid in Mac OS X Server version 10.5 includes a controller for computational grids and an agent that allows the server's processor to work on jobs submitted to a grid. The agent is also available in computers using Mac OS X v10.3 or v10.4.

What's New in Xgrid Administration

Xgrid service, Xgrid Admin, and high performance computing (HPC) in Mac OS X Server v10.5 Leopard include the following valuable new features.

- Improved security with Xgrid superuser access controls
- New Xgrid service configuration assistant
- Logging improvements

What's in This Guide

This guide is organized as follows:

- Part I—Xgrid Administration. The chapters in this part of the guide introduce you to Xgrid service and the applications and tools available for administering xgrid.
- Part II—Configuring High Performance Computing. The chapters in this part of the guide introduce you to HPC and the applications and tools available for administering HPC.

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 This Guide

The following list contains suggestions for using this guide:

- Read the guide in its entirety. Subsequent sections might build on information and recommendations discussed in prior sections.
- The instructions in this guide should always be tested in a nonoperational environment before deployment. This nonoperational environment should simulate, as much as possible, the environment where the computer will be deployed.

Using Onscreen Help

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

To get help for an advanced configuration of Leopard Server:

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

The help for Server Admin and Workgroup Manager contains instructions taken from *Server Administration* and other advanced administration guides described in "Advanced Server Administration Guides," next.

To see the latest server help topics:

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

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

Advanced Server Administration Guides

Getting Started covers basic installation and initial setup methods for a standard, workgroup, or advanced configuration of Leopard Server. An advanced guide, *Server Administration*, covers advanced planning, installation, setup, and more. A suite of additional guides, listed below, covers advanced planning, setup, and management of individual services. You can get these guides in PDF format from the Mac OS X Server documentation website at www.apple.com/server/documentation.

This guide ...	tells you how to:
<i>Getting Started and Mac OS X Server Worksheet</i>	Install Mac OS X Server and set it up for the first time.
<i>Command-Line Administration</i>	Install, set up, and manage Mac OS X Server using UNIX command-line tools and configuration files.
<i>File Services Administration</i>	Share selected server volumes or folders among server clients using the AFP, NFS, FTP, and SMB/CIFS protocols.
<i>iCal Service Administration</i>	Set up and manage iCal shared calendar service.
<i>iChat Service Administration</i>	Set up and manage iChat instant messaging service.
<i>Mac OS X Security Configuration</i>	Make Mac OS X computers (clients) more secure, as required by enterprise and government customers.
<i>Mac OS X Server Security Configuration</i>	Make Mac OS X Server and the computer it's installed on more secure, as required by enterprise and government customers.
<i>Mail Service Administration</i>	Set up and manage IMAP, POP, and SMTP mail services on the server.
<i>Network Services Administration</i>	Set up, configure, and administer DHCP, DNS, VPN, NTP, IP firewall, NAT, and RADIUS services on the server.
<i>Open Directory Administration</i>	Set up and manage directory and authentication services, and configure clients to access directory services.
<i>Podcast Producer Administration</i>	Set up and manage Podcast Producer service to record, process, and distribute podcasts.
<i>Print Service Administration</i>	Host shared printers and manage their associated queues and print jobs.
<i>QuickTime Streaming and Broadcasting Administration</i>	Capture and encode QuickTime content. Set up and manage QuickTime streaming service to deliver media streams live or on demand.
<i>Server Administration</i>	Perform advanced installation and setup of server software, and manage options that apply to multiple services or to the server as a whole.
<i>System Imaging and Software Update Administration</i>	Use NetBoot, NetInstall, and Software Update to automate the management of operating system and other software used by client computers.
<i>Upgrading and Migrating</i>	Use data and service settings from an earlier version of Mac OS X Server or Windows NT.
<i>User Management</i>	Create and manage user accounts, groups, and computers. Set up managed preferences for Mac OS X clients.

This guide ...	tells you how to:
<i>Web Technologies Administration</i>	Set up and manage web technologies, including web, blog, webmail, wiki, MySQL, PHP, Ruby on Rails, and WebDAV.
<i>Xgrid Administration and High Performance Computing</i>	Set up and manage computational clusters of Xserve systems and Mac computers.
<i>Mac OS X Server Glossary</i>	Learn about terms used for server and storage products.

Viewing PDF Guides on Screen

While reading the PDF version of a guide on screen:

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

Printing PDF Guides

If you want to print a guide:

- Save ink or toner by not printing the cover page.
- Save color ink on a color printer by looking in the panes of the Print dialog for an option to print in grays or black and white.
- Maximize the printed page image by changing the Scale setting in the Page Setup dialog. Try 122% with Paper Size set to US Letter. (PDF pages are 7.5 by 9 inches except *Getting Started*, which is CD size, 125 by 125 mm.)
- Reduce the bulk of the printed document and save paper by printing more than one page per sheet of paper. In the Print dialog, choose Layout from the untitled pop-up menu. If your printer supports two-sided (duplex) printing, select one of the Two-Sided options. Otherwise, choose 2 from the Pages per Sheet pop-up menu, and optionally choose Single Hairline from the Border menu.

Getting Documentation Updates

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

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

Getting Additional Information

For more information, consult these resources:

- *Read Me documents*—important updates and special information. Look for them on the server discs.
- *Mac OS X Server website* (www.apple.com/macosx/server)—gateway to extensive product and technology information.
- *Apple Service & Support website* (www.apple.com/support)—access to hundreds of articles from Apple’s support organization.
- *Apple customer training* (train.apple.com)—instructor-led and self-paced courses for honing your server administration skills.
- *Apple discussion groups* (discussions.info.apple.com)—a way to share questions, knowledge, and advice with other administrators.
- *Apple mailing list directory* (www.lists.apple.com)—subscribe to mailing lists so you can communicate with other administrators using email.
- *Open Source website* (developer.apple.com/darwin/)—Access to Darwin open source code, developer information, and FAQs.

Part I: Xgrid Administration

Use the chapters in this part of the guide to learn about Xgrid service and the applications and tools available for administering Xgrid.

- Chapter 1 **Introducing Xgrid Service**
- Chapter 2 **Setting Up and Configuring Xgrid Service**
- Chapter 3 **Managing a Grid**
- Chapter 4 **Planning and Submitting Xgrid Jobs**
- Chapter 5 **Solving Xgrid Problems**

Use this chapter to learn about what Xgrid is and how it can help you.

You use Xgrid to create grids of multiple computers and distribute complex jobs among them for high-throughput computing.

Xgrid, a technology in Mac OS X Server and Mac OS X, simplifies deployment and management of computational grids. Xgrid enables administrators to group computers in grids or clusters, and enables users to easily submit complex computations to groups of computers (local, remote, or both), as either an ad hoc grid or a centrally managed cluster.

About Xgrid and Computational Grids

Xgrid makes it easy to turn an ad hoc group of Mac systems into a low-cost supercomputer. Xgrid is ideal for individual researchers, specialized collaborators, and application developers. For example:

- Scientists can search biological databases on a cluster of Xserve systems.
- Engineers can perform finite element analyses on their workgroup's desktops.
- Animators can render images using Mac systems across multiple corporate locations.
- Research teams can enlist colleagues and interested laypeople in Internet-scale volunteer grids to perform long-running scientific calculations.
- Anyone needing to perform CPU-intensive calculations can simultaneously run a single job across multiple computers, dramatically improving throughput and responsiveness.

With Xgrid functionality integrated into Mac OS X Server, system administrators can quickly enable Xgrid on Mac systems throughout their company, turning idle CPU cycles into a productive cluster at no incremental cost.

How Xgrid Works

Xgrid creates multiple tasks for each job and distributes those tasks among multiple nodes. These nodes can be desktop computers running Mac OS X v10.3 or later, or server computers running Mac OS X Server v10.4 or later.

Many desktop computers sit idle during the day, in evenings, and on weekends. The assembly of these systems into a computational grid is known as *desktop recovery*. This method of grid construction enables you to vastly improve your computational capacity without purchasing extra hardware, and Xgrid makes the software configuration a straightforward task.

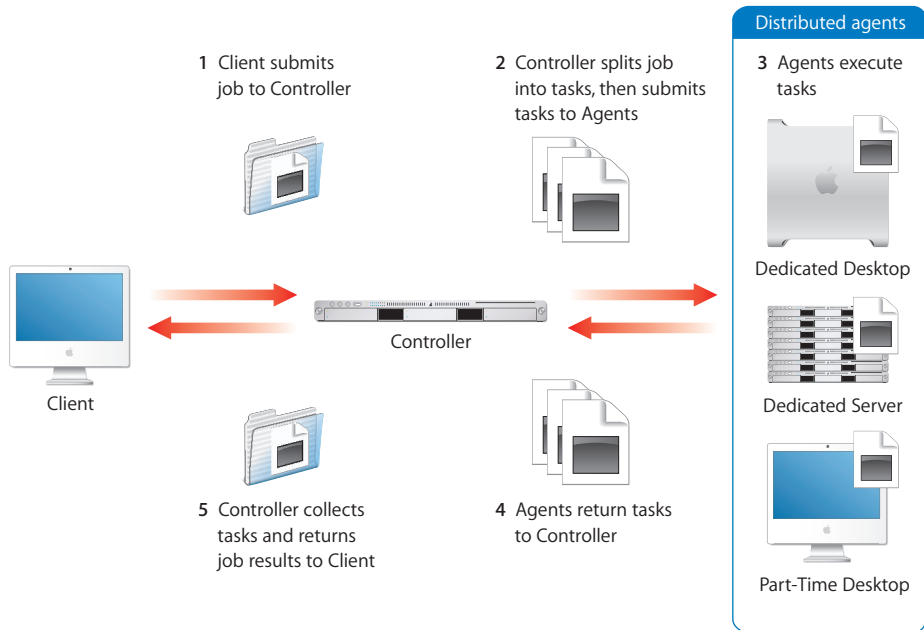
For a server to function as a controller, Xgrid requires Mac OS X Server v10.4 or later, with a minimum of 256 MB of RAM. To operate as an agent in a grid, Xgrid requires Mac OS X v10.3 or later, with a minimum of 128 MB of RAM (256 MB advisable). All Xgrid participants must have a network connection. As always, the more RAM a system has, the better it performs, particularly for high-performance computing applications.

A *grid* is a group of computers working together to solve a single problem. The systems in a grid can be loosely coupled, geographically dispersed and, to some extent, heterogeneous. In contrast, systems in a *cluster* are often homogeneous, collocated, and strictly managed.

Highly dispersed grids, such as SETI@Home, enable individuals to donate their spare processor cycles to a cause. In office environments, large rendering or simulation jobs can be distributed across all the systems left idle overnight. These can even be used to augment a dedicated computational cluster, which is available to Xgrid clients at all times.

These distinct grid configurations are explained in “Common Types of Grids and Grid Computing Styles” on page 20.

The illustration below gives an example of how a grid handles a job.



Xgrid has no limitations on the amount of computational power it can support. The performance of the grid depend on the systems participating, the software running, and the network, among other factors. However, individual applications strongly influence the performance of the grid.

You determine if an application is improved by being deployed on a computational grid. In the best case, application performance may scale linearly with the size of the grid. In the worst case, the addition of agents to a grid can cause a job to complete in even more time than if there were fewer agents. (In such a situation, tasks become so small that the overhead associated with distributing the increased number of tasks supersedes the performance gain of using more agents.) You should be aware of these considerations.

Many proprietary projects enable you to participate in a large computational grid. Often these projects, such as SETI@Home and FightAIDS@Home, are tied to a specific scientific purpose. They usually have easy-to-install software that enables any volunteer to participate in that particular project, and they frequently take the form of a screen saver or background process.

You don't need to think in terms of thousands or millions of seldom-used computers to see the significance of a computational grid. For example, computers used by university students or corporate employees often work fewer hours than the hours they sit idle at night or on weekends. These computers could contribute productively to the work of a grid without diminishing their usefulness to the students or employees.

Other grid projects are designed for large-scale computational grids, such as the Globus Alliance (a group founded by universities and researchers), with flexible resource management tools and more intelligent grid deployment methods. Instead of developing neatly packaged applications for a specific grid, such projects provide comprehensive frameworks for application deployment.

Xgrid enables users to participate in a computational grid of their choice while still providing the flexibility of a more generic framework for grid developers when deploying grid applications. Xgrid provides the primary benefits of both.

The advantages of the Xgrid technology include:

- Easy grid configuration and deployment
- Straightforward yet flexible job submission
- Automatic controller discovery by agents and clients
- Flexible architecture based on open standards
- Support for the UNIX security model, including Kerberos single sign-on or regular password authentication
- Choice between a command-line interface or an API-based model for grid interaction

Common Types of Grids and Grid Computing Styles

Xgrid can be used in tightly coupled clusters, worldwide grids, and everything in between. This immense flexibility enables you to deploy grids of almost any nature. Three main topologies are commonly used for Xgrid deployments, discussed as follows:

- "Xgrid Clusters" on page 20
- "Local Grids" on page 21
- "Distributed Grids" on page 21

Xgrid Clusters

Computational clusters are sets of systems dedicated to computation. In a cluster, systems are typically co-located in a rack, connected using gigabit Ethernet or another high-performance network, and strictly managed for maximum performance.

Cluster systems are often entirely homogeneous: their operating systems are the same versions, they have the same software installed, and they generally have the same processor, disk, and RAM configurations.

Xgrid enables administrators to easily configure the distributed resource management functionality of the cluster. Each server in the system runs the agent software, and the head node in the cluster runs the controller software.

Xgrid distributes tasks across the cluster. In clusters, failure rates are generally very low. Systems are rarely, if ever, offline, and their resources are not shared with general user tasks. Clusters are the most efficient but most expensive model of distributed computing.

Local Grids

Systems that are under common administration in a company, university computer lab, or other managed environment can often be easily assembled into a grid for desktop recovery. These systems are often on a local area network (LAN) and they are generally managed by a single organization. As a result, they provide good network performance and offer substantial manageability.

Because these systems are often also used as day-to-day workstations, users can easily interrupt grid tasks by moving the mouse, resetting the system, or even accidentally disconnecting the system from the network. In such cases, a task might fail as part of an Xgrid job the Xgrid controller eventually reassigns the failed task to another agent, and the job completes successfully.

In local grids, performance is limited by such situations and by the varying performance of any given agent on the grid.

Distributed Grids

When a system is permitted to donate its time, a distributed grid is formed.

The Xgrid agent enables a user to specify any IP address or host name for its controller. By specifying a grid, a user can dedicate his or her CPU time to that grid no matter where the controller is located.

The manager of the controller has no direct management control or knowledge of the agent system but is nonetheless able to harness its CPU time.

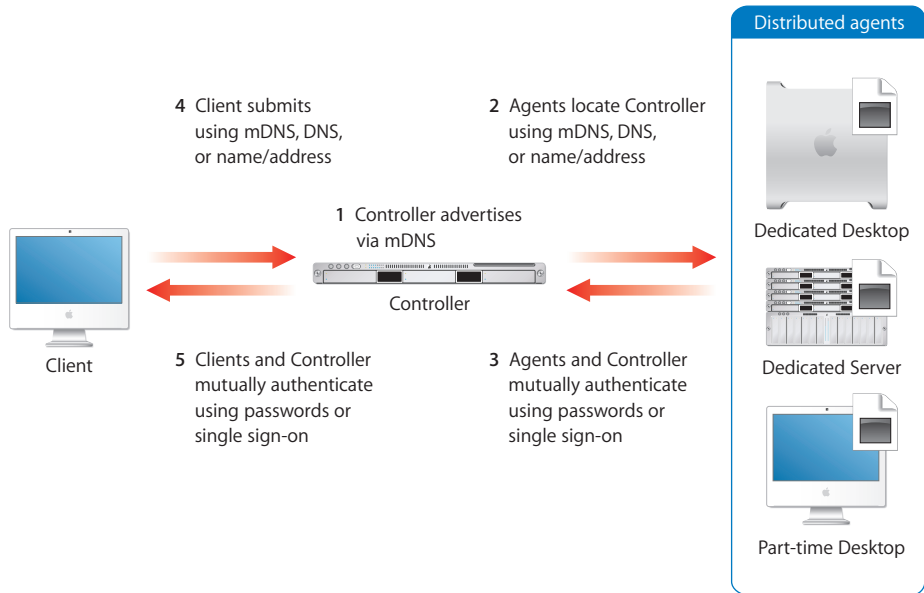
Distributed grids have very high failure rates for jobs but place a very low burden for the grid administrator. With very, very large jobs, high task failure rates may not substantially affect the performance of the grid if such failures can be rapidly reassigned to other available agents.

Network performance can also be a consideration because data is sent over the Internet, rather than over a local network, to agents connected to a grid. The monetary cost of such distributed grids is extremely low.

Xgrid Components

The Xgrid three-tier architecture simplifies the distribution of complicated tasks. Its user clients, grid controllers, and computational agents work together to streamline the process of assembling nodes, submitting jobs, and retrieving results.

The illustration below gives an example of the Xgrid components and the process of auto configuration for a grid.



The primary components of a computational grid perform the following functions:

- An agent runs one task at a time per CPU; a dual-processor computer can run two tasks simultaneously.
- A controller queues tasks, distributes those tasks to agents, and handles task reassignment.
- A client submits jobs to the Xgrid controller in the form of multiple tasks. (A client can be any computer running Mac OS X v10.4 or later or Mac OS X Server v10.4 or later.)

In principle, the agent, controller, and client can run on the same server, but it is often more efficient to have a dedicated controller node.

Agent

Xgrid agents run the computational tasks of a job. In Mac OS X Server, the agent is turned off by default. When an agent is turned on and becomes active at startup, it registers with a controller. (An agent can be connected to only one controller at a time.) The controller sends instructions and data to the agent as needed for the controller's jobs. After it receives instructions from the controller, the agent performs its assigned tasks and sends the results back to the controller.

By default, agents seek to bind to the first available controller on the LAN. Alternatively, you can specify that it bind to a specific controller.

You can also specify whether an agent is always available or is available only when the computer is idle. A computer is considered idle when it has no mouse or keyboard input and ignores CPU and network activity. If a user returns to a computer that is running a grid task, the computer continues to run the task until it is finished.

By default, the agent on Mac OS X Server is dedicated and the agent on a Mac OS X computer (not a server) is configured to accept tasks only when the computer has had no user input for 15 minutes.

For details about configuring an agent, see “Configuring an Xgrid Agent (Mac OS X Server)” on page 32.

For information about managing agents, see “Managing Agents” on page 42.

Client

Any system can be an Xgrid client if it is running Mac OS X v10.4 or later and has a network connection to the Xgrid controller system. In general, the client can connect to only a single controller.

Depending on how a controller is configured, the client must supply a password or be authenticated by Kerberos (single sign-on) before submitting a job to the grid.

A user submits a job to the controller from a system running the Xgrid client software, usually a command-line tool accessed with the Terminal application. The job can specify the controller or use multicast DNS (mDNS) to dynamically discover the first available controller. When the job is complete, the controller notifies the client and the client can retrieve the results of the job.

For information about client authentication to the controller, see “Setting Up Grid Authentication” on page 34.

Controller

The Xgrid controller manages the communications among the computational resources of a grid. The controller requires Mac OS X Server v10.4 or later. The controller accepts network connections from clients and agents. It receives job submissions from clients, divides the jobs into tasks, dispatches tasks to agents, and returns results to the clients.

Although there can be more than one Xgrid controller running on a subnet, there can only be one controller per logical grid. Each controller can have an arbitrary number of agents connected, but Apple has tested 128 agents per controller.

However, there is no software limitation on the number of agents, and users of Xgrid can choose to exceed 128 agents on a controller at their own risk, with a theoretical maximum equal to the number of available sockets on the controller system.

For details about setting up an Xgrid controller, see “Configuring Controller Settings” on page 30.

For information about managing controllers and grids, see “Managing the Xgrid Controller” on page 40.

Jobs

A job is a collection of execution instructions that can include data and executables. Xgrid can run scripts, utilities, and custom software (anything that doesn't require user interaction).

A client submits a job to the grid. The controller accepts the job and its associated files, divides the job into tasks, and then distributes the tasks to agents. Agents accept the tasks, perform the calculations, and return the results to the controller, which aggregates them and returns them to the clients.

For more information about jobs, see “Structuring Jobs for Xgrid” on page 47 and “Submitting a Job” on page 48.

Requirements and Capacities

Xgrid is designed to scale from small clusters of a few computers up to large organization-wide grids. Xgrid supports up to 128 agents, any number of jobs comprising up to 100,000 queued tasks, up to 128 MB of submitted data per job, and up to 128 MB of results per job. These are recommended limits and are not enforced by the software. You may choose to exceed these limits at your own risk.

Use this chapter to plan your grid and set up the Xgrid agent and controller.

Xgrid simplifies deployment and management of computational grids. Using Server Admin you can configure Xgrid to set up computer groups (grids or clusters) and allow users to easily submit complex computations to these grids (local, remote, or both), as either an ad hoc grid or a centrally managed cluster.

Setup Overview

Here is an overview of the steps for setting up Xgrid service:

Step 1: Before you begin

See “Before Setting Up Xgrid Service” on page 26. Identify the Xgrid environment you need. Before configuring Xgrid, you must make some decisions about the grid.

Step 2: Turn Xgrid service on

Prior to configuring, turn on Xgrid service. See “Turning Xgrid Service On” on page 28.

Step 3: (Optional) Use the Xgrid service configuration assistant to configure Xgrid

If you choose to, you can configure Xgrid using the Xgrid service configuration assistant. This assistant helps with Xgrid configuration by automating many of the settings you make. See “Configuring Xgrid with the Xgrid Service Configuration Assistant” on page 28.

Step 4: Configure Xgrid controller settings

Configure your server as an Xgrid controller using Server Admin. See “Configuring Controller Settings” on page 30.

Step 5: Start Xgrid service

Start Xgrid service on the server using Server Admin. See “Starting Xgrid Service” on page 31.

Step 6: Configure Xgrid agent settings (Mac OS X Server)

Configure your server as an Xgrid agent using Server Admin. See “Configuring an Xgrid Agent (Mac OS X Server)” on page 32.

Step 7: Configuring Xgrid agent settings (Mac OS X)

Configure computers as Xgrid agents by using Sharing Preferences. See “Configuring an Xgrid Agent (Mac OS X)” on page 33.

Before Setting Up Xgrid Service

Before configuring Xgrid service, you must define the grid environment you’ll create. In particular, you must decide the following:

- The kind of authentication to use. See “Authentication Methods for Xgrid” on page 26.
- Where to host your controller. See “Hosting the Grid Controller” on page 28.
- How you will manage the controller. See “Managing Xgrid Service” on page 37 and “Monitoring Grid Activity” on page 46.

Authentication Methods for Xgrid

You can configure Xgrid with or without authentication. If you choose to require authentication of controllers to mutually authenticate with clients and agents, you can choose Single Sign-On or Password-Based Authentication. The following authentication options are available:

- Single Sign-On
- Password-Based Authentication
- No Authentication

You set up an Xgrid controller using Server Admin. You can specify the type of authentication for agents and clients. The passwords entered in Server Admin for the controller must match those entered for each agent and client.

Consider these points when establishing passwords for agents and clients:

- **Kerberos authentication (single sign-on or SSO).** If you use Kerberos authentication for agents or clients, the server that’s the Xgrid controller must be configured for Kerberos, in the same realm as the server running the Kerberos domain controller (KDC) system, and bound to the Open Directory master.
The agent uses the host principal found in the `/etc/krb5.keytab` file. The controller uses the Xgrid service principal found in the `/etc/krb5.keytab` file.
- **Agents.** The agent determines the authentication method. The controller must conform to that method and password (if a password is used). When an agent is configured with a standard password (not SSO), you must use the same password for agents when you configure the controller. If the agent has specified SSO, the correct service principal and host principals must be available.

- **Clients.** If your server is the controller for a grid, be sure that Mac OS X and Mac OS X Server clients use the correct authentication method for the controller.
A client cannot submit a job to the controller unless the user chooses the correct authentication method and enters their password correctly, or has the correct ticket-granting ticket from Kerberos.

For more information, see “Setting Up Grid Authentication” on page 34.

Single Sign-On (SSO)

SSO is the most powerful and flexible form of authentication. It leverages the Open Directory and Kerberos infrastructures in Mac OS X Server to manage authentication behind the scenes, without user intervention.

Each Xgrid participant must have a Kerberos principal. The clients and agents obtain ticket-granting tickets for their principal, which is used to obtain a service ticket for the controller service principal. The controller looks at the ticket granted to the client to determine the user’s principal and verifies it with the relevant service access control lists (SACLs) and groups to determine privileges.

Generally, you should use this option if any of the following conditions are true:

- You already have SSO in your environment.
- You have administrator control over all agents and clients in use.
- Jobs must run with special privileges (such as for local, network, or SAN file system access).

Password-Based Authentication

When you can’t use SSO, you can require password authentication. You may not be able to use SSO if:

- Potential Xgrid clients are not trusted by your SSO domain (or you don’t have one)
- You want to use agents across the Internet or that are outside your control
- It is an ad hoc grid, without the ability to prearrange a web of trust

In these situations, your best option is to specify a password. You have two distinct password settings: one for controller-client and one for controller-agent. For security reasons these should be different passwords.

Note: You can also create hybrid environments, such as with client-controller authentication done using passwords but controller-agent authentication done using SSO (or vice versa).

No Authentication

This option is suitable only for testing a private network in a home or a lab that is inaccessible from any untrusted computer, or when none of the jobs or the computers contain sensitive or important information.

Otherwise, do not use this option. It creates a potential security hole (because anyone can connect or run a job) and should never be used on a system exposed to the Internet, especially when potentially sensitive data is involved.

If you choose to use no authentication, agents can join the grid and clients can submit jobs to the grid without authenticating.

Hosting the Grid Controller

The primary requirement for a controller is that it must be network-accessible to clients and agents. In some cases this may mean the controller must be placed outside an organizational firewall (or inside a buffer zone); otherwise, you would need to open up port 4111 so the controller can be contacted.

It is much simpler (though not essential) for the controller to be on the same subnet as the agents and usual clients, so they can discover each other using Bonjour. If that's not feasible, host the controller on a server with a fixed IP address and fully qualified DNS name (or alternatively, using Dynamic DNS and a service lookup entry) so that agents and clients know where to find it.

Turning Xgrid Service On

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

To turn Xgrid service on:

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

Configuring Xgrid with the Xgrid Service Configuration Assistant

You can set up Xgrid service by configuring the controller and agent using the Xgrid service configuration assistant. This optional configuration assistant guides you through setting up a server to host a grid or join an existing grid.

Before this assistant proceeds, your server must have access to a directory server that provides Kerberos services.

Configuring Xgrid to Host a Grid Using the Xgrid Service Configuration Assistant

Use the Xgrid service configuration assistant to configure the Xgrid agent and controller to run on this server. This also configures a network file system.

To set up Xgrid to host a grid using the Xgrid service configuration assistant:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 In the expanded Servers list, click Xgrid.
- 4 Click Overview.
- 5 Click Configure Xgrid Service (at the lower right).
This opens the Xgrid service configuration assistant.
- 6 Click Continue.
- 7 Choose “Host a grid,” then click Continue.
- 8 Enter the username and password for the directory administrator to authenticate with the directory domain displayed, then click Continue.
- 9 Review and confirm your configuration settings, then click Continue.
This restarts Xgrid service using your settings.
- 10 Click Close.

Configuring Xgrid to Join a Grid Using Xgrid Service Configuration Assistant

Use the Xgrid service configuration assistant to configure the Xgrid agent to run on this server. Joining a grid means that an agent is set up on this server and is bound to an existing controller.

To join a grid using the Xgrid service configuration assistant:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 In the expanded Servers list, click Xgrid.
- 4 Click Overview.
- 5 Click Configure Xgrid Service (at the lower right).
This opens the Xgrid service configuration assistant.
- 6 Click Continue.
- 7 Choose “Join a grid,” then click Continue.

- 8 Specify the controller you want to bind your agent to.
Select “Browse Bonjour-discoverable controllers” to view and select from available controllers.
Select “Use controller with hostname” to enter the hostname of a specific controller.
- 9 Click Continue.
- 10 Review and confirm your configuration settings, then click Continue.
This restarts Xgrid service using your settings.
- 11 Click Close.

Setting Up Xgrid Service

You set up Xgrid service by configuring two groups of settings on the Settings pane for Xgrid service in Server Admin:

- **Controller.** Use to configure your server as an Xgrid controller and set client and agent authentication.
- **Agent.** Use to configure your server as an Xgrid agent, to specify the controller, and to set controller authentication.

The following section describes how to configure these settings. An additional section tells you how to start Xgrid service when you finish. (By default, the Xgrid controller and agent are disabled.)

Important: If you specify a password, the agent and controller must use the same password or must authenticate using Kerberos (SSO). For information about authentication options, see “Setting Passwords for Xgrid” on page 34.

Xgrid and Multiple Network Interfaces

On a server with multiple network interfaces, Mac OS X Server makes Xgrid service available over all interfaces. You can’t configure Xgrid service separately for each interface.

Configuring Controller Settings

You use Server Admin to configure an Xgrid controller. When configuring the controller, you can also set a password for any agent using the grid and for any client that submits a job to the grid.

To configure an Xgrid controller:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 In the expanded Servers list, click Xgrid.

- 4 Click Settings.
- 5 Click Controller.
- 6 Click “Enable controller service.”
- 7 From the Client Authentication pop-up menu, choose one of the following authentication options for clients and enter the password.
 - **Password** requires that the agent and controller use the same password.
 - **Kerberos** uses SSO authentication for the agent’s administrator.
 - **None** does not require a password for the agent. This option provides no protection from potentially malicious use of your grid. With no authentication, a malicious agent could receive tasks and potentially access sensitive data.

For details about password options, see “Setting Up Grid Authentication” on page 34.

- 8 From the Agent Authentication pop-up menu, choose from the following authentication options for agents and enter the password.
 - **Password** requires that the agent and controller use the same password.
 - **Kerberos** uses SSO authentication for the agent’s administrator.
 - **Any** uses any authentication available for the agent’s administrator.
 - **None** does not require a password for the agent. This option provides no protection from potentially malicious use of your grid. With no authentication, a malicious agent could receive tasks and potentially access sensitive data.

For information about password options, see “Setting Up Grid Authentication” on page 34.

- 9 Click Save.

Important: If you require authentication, the agent and controller must use the same password or must authenticate using Kerberos (SSO). For information about authentication options, see “Setting Up Grid Authentication” on page 34.

Starting Xgrid Service

Use Server Admin to start Xgrid service.

The Xgrid service must be running for your server to control a grid or participate in a grid as an agent.

For details about using the server as an agent and controller, see “Configuring an Xgrid Agent (Mac OS X Server)” on page 32 and “Configuring Controller Settings” on page 30.

After you start Xgrid, it restarts when the server is restarted.

To start Xgrid service:

- 1 Open Server Admin and connect to the server.

- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 In the expanded Servers list, click Xgrid.
- 4 Click the Start Xgrid button (below the Servers list).

Configuring an Xgrid Agent (Mac OS X Server)

You use Server Admin to set up your server as an Xgrid agent. In addition, you can associate the agent with a specific controller or permit it to join a grid, specify when the agent accepts tasks, and set a password that the controller must recognize.

To configure an Xgrid agent on the server:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 In the expanded Servers list, click Xgrid.
- 4 Click Settings.
- 5 Click Agent.
- 6 Click “Enable agent service.”
- 7 Specify a controller by choosing its name in the Controller pop-up menu or by entering the controller name.

By default, the agent uses the first available controller.

Note: An agent can find a controller in one of three ways: a specific hostname or IP address, the first available controller that advertises on Bonjour on the local subnet, or to a specific Bonjour service name.

- 8 Specify when the agent will accept tasks.

Tasks can be accepted when the computer is idle or always.

A computer is considered idle when it has no mouse or keyboard input and ignores CPU and network activity. If a user returns to a computer that is running a grid task, the computer continues to run the task until it is finished.

- 9 From the pop-up menu, choose one of the following authentication options and enter the password.

For details, see “Setting Up Grid Authentication” on page 34.

- **Password** requires that the agent and controller use the same password.
- **Kerberos** uses SSO authentication for the agent’s administrator.
- **None** does not require a password for the agent. This option provides no protection from potentially malicious use of your grid. With no authentication, a malicious agent could receive tasks and potentially access sensitive data.

- 10 Click Save.

Important: If you require authentication, the agent and controller must use the same password or must authenticate using Kerberos SSO. For details about authentication option, see “Setting Up Grid Authentication” on page 34.

Configuring an Xgrid Agent (Mac OS X)

You use Sharing preferences to set up client computers as Xgrid agents. In addition, you can associate the agent with a specific controller or permit it to join any grid, specify when the agent accepts tasks, and set a password that the controller must recognize.

To configure an Xgrid agent on a client:

- 1 On the client computer, open Sharing preferences and click Services.
- 2 Click Xgrid and then click Configure.
- 3 Specify a controller by choosing its name in the Controller pop-up menu or by entering the controller name.

By default, the agent uses the first available controller.

Note: An agent can find a controller in one of three ways: a specific hostname or IP address, the first available controller that advertises on Bonjour on the local subnet, or to a specific Bonjour service name.

- 4 Specify when the agent will accept tasks.

Tasks can be accepted when the computer is idle or always.

A computer is considered idle when it has no mouse or keyboard input and ignores CPU and network activity. If a user returns to a computer that is running a grid task, the computer continues to run the task until it is finished.

- 5 Choose one of the following authentication options from the pop-up menu and enter the password.

For more information, see “Setting Up Grid Authentication” on page 34.

- **Password** requires that the agent and controller use the same password.
- **Kerberos** uses SSO authentication for the agent’s administrator.
- **None** does not require a password for the agent. This option provides no protection from potentially malicious use of your grid. With no authentication, a malicious agent could receive tasks and potentially access sensitive data.

- 6 Click OK.

Important: If you require authentication, the agent and controller must use the same password or must authenticate using Kerberos (SSO). For more information about authentication options, see “Setting Up Grid Authentication” on page 34.

- 7 Click Start to turn Xgrid sharing on.

Setting Up Grid Authentication

You can configure Xgrid to require authentication of controllers, clients, and agents. For more information, see “Authentication Methods for Xgrid” on page 26.

Setting Up Kerberos for Xgrid

You use Server Admin to configure Kerberos as the authentication method for your Xgrid. Kerberos authentication uses SSO.

To configure Kerberos authentication:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 In the expanded Servers list, click Xgrid.
- 4 Click Settings.
- 5 Click Agent.
- 6 Click “Enable agent service.”
- 7 For the authentication option for the agent, choose Kerberos from the Controller Authentication pop-up menu.
- 8 Click Controller.
- 9 Click “Enable controller service.”
- 10 For the authentication option for the client, choose Kerberos from the Client Authentication pop-up menu.
- 11 For the authentication option for the agent, choose Kerberos from the Agent Authentication pop-up menu.
- 12 Click Save and restart the service.

Setting Passwords for Xgrid

You use Server Admin to configure your Xgrid controllers to authenticate clients and agents using password authentication. Password authentication requires that the agent and controller use the same password.

You specify password options in Server Admin as part of configuring the agent and controller. See “Configuring an Xgrid Agent (Mac OS X Server)” on page 32 and “Configuring Controller Settings” on page 30.

To configure password authentication:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.

- 3 In the expanded Servers list, click Xgrid.
- 4 Click Settings.
- 5 Click Agent.
- 6 Click "Enable agent service."
- 7 For the authentication option for the agent, choose Password from the Controller Authentication pop-up menu and enter a password.
- 8 Click Controller.
- 9 Click "Enable controller service."
- 10 For the authentication option for the client, choose Password from the Client Authentication pop-up menu and enter a password.
- 11 For the authentication option for the agent, choose Password from the Agent Authentication pop-up menu and enter a password.

You can also choose Any from the Agent Authentication pop-up menu to permit any method of authentication.

Note: Password authentication requires that the agent and controller use the same password.
- 12 Click Save and restart the service.

Managing Client Access

Server Admin in Mac OS X Server enables you to configure service access control lists (SACLs), which enable you to specify which users and groups have access to Xgrid and which administrators can manage it.

Using SACLs enables you to add another layer of access control in addition to password and Kerberos authentication. Only users and groups listed in an SACL have access to its corresponding service.

Setting SACL Permissions for Users and Groups

You use Server Admin to set SACL permissions for users and groups to access Xgrid service.

To set user and group SACL permissions for Xgrid service:

- 1 Open Server Admin and connect to the server.
- 2 Click Settings.
- 3 Click Access.
- 4 Click Services.
- 5 Select the level of restriction you want for the services:

To restrict access to all services, select “For all services.”

To set access permissions for individual services, select “For selected services below,” then select a service from the Service list.

- 6 To provide unrestricted access to services, click “Allow all users and groups.”
- 7 To restrict access to users and groups:
 - a Select “Allow only users and groups below.”
 - b Click the Add (+) button to open the Users and Groups drawer.
 - c Drag users and groups from the Users and Groups drawer to the list.
- 8 Click Save.

Setting SACL Permissions for Administrators

Use Server Admin to set SACL permissions for administrators to monitor and manage Xgrid service.

To set administrator SACL permissions for Xgrid service:

- 1 Open Server Admin and connect to the server.
- 2 Click Settings.
- 3 Click Access.
- 4 Click Administrators.
- 5 Select the level of restriction you want for the services:

To restrict access to all services, select “For all services.”

To set access permissions for individual services, select “For selected services below,” then select a service from the Service list.
- 6 Open the Users and Groups drawer by clicking the Add (+) button.
- 7 From the Users and Groups drawer, drag users and groups to the list.
- 8 Set user permissions:

To grant administrator access, choose Administer from the Permission pop-up menu next to the user name.

To grant monitoring access, choose Monitor from the Permission pop-up menu next to the user name.
- 9 Click Save.

Managing Xgrid Service

This section describes typical day-to-day tasks you might perform after you set up Xgrid service on your server. For information about initial setup, see “Setting Up Xgrid Service” on page 30.

You can monitor and manage grids using Xgrid Admin. For more information, see Chapter 3, “Managing a Grid.”

Viewing Xgrid Service Status

You can use Server Admin to view the status of Xgrid service.

To view Xgrid service status:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 From the expanded Servers list, select Xgrid.
- 4 Click Overview to see whether the service is running, when it started, agent and controller information, the number of jobs running and pending, and the amount of processor power available and used.
- 5 Click Logs to review the system, controller, and agent logs.
Use the View pop-up menu to choose which log to view.

Viewing Xgrid Service Logs

You can use Server Admin to view the Xgrid system, controller, and agent logs for Xgrid service.

To view logs:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 From the expanded Servers list, select Xgrid.
- 4 Click Logs, then use the Show pop-up menu to choose System Log (Xgrid), Xgrid Controller Log, or Xgrid Agent Log.

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

From the Command Line

You can also view the Xgrid service log at `/var/log/system.log` using the `cat` or `tail` commands in Terminal.

Stopping Xgrid Service

You use Server Admin to stop Xgrid service.

To stop Xgrid service:

- 1 Open Server Admin and connect to the server.
- 2 Click the triangle to the left of the server.
The list of services appears.
- 3 From the expanded Servers list, select Xgrid.
- 4 Click the Stop Xgrid button (below the Servers list).

From the Command Line

You can also stop Xgrid service immediately by using the `serveradmin` command in Terminal.

Use this chapter to learn how to use the Xgrid Admin application to manage grids, add controllers and agents, and work with jobs.

After you set up an Xgrid controller, you can use Xgrid Admin to manage a grid. You can use Xgrid Admin on the server or on a remote computer that is running Mac OS X v10.4 or later.

You can manage one or more computational grids with Xgrid Admin. A computational grid is a fixed group of agents with a dedicated queue. There can be multiple grids per controller but an agent can belong to only one grid. You cannot move an agent between grids while a job (or a task) is running.

Using Xgrid Admin

Xgrid Admin is a tool you use to monitor one or more grids and manage agents and jobs.

With Xgrid Admin, you can:

- Check the status of a grid and its activity, including the number of agents working and available, processing power in use and available, and the number of jobs running and pending
- Add or remove controllers and grids to manage
- See a list of agents in a grid and the CPU power available and in use for each agent
- Add or remove agents in a grid
- See a list of jobs in a grid, the date and time each job was submitted, its progress, and the active CPU power for the job
- Remove jobs in a grid
- Stop a job in progress
- Restart a job that was stopped or is complete

Xgrid Admin provides controls in its graphical interface and menu commands for all of its options.

Note: You can also use the Xgrid command-line tool to perform these tasks. For more information about using the command-line tool, see Chapter 4, “Planning and Submitting Xgrid Jobs.”

Status Indicators in Xgrid Admin

Xgrid Admin provides status indicators, which are small color bubbles indicating the status of controllers, agents, and jobs. The color indicators are:

- Colorless = controller or agent is offline, job is pending
- Gray = job is submitting
- Green = controller is connected, agent is working, job is running
- Yellow = agent is available but not running
- Red = agent is unavailable, job is failed or canceled
- Blue = job is complete

Managing the Xgrid Controller

In general, you manage the Xgrid controller like any other service running on Mac OS X Server, using Server Admin to manage which processes are running and using Xgrid Admin to manage the agent and job queues on the controller.

The amount of management required also depends on how many queues you have and the number (and temperament) of the users who submit jobs.

Xgrid uses a simple first-in, first-out (FIFO) queue for scheduling each grid, which means that as the administrator you must obtain your colleagues’ cooperation to make sure resources are allocated correctly among multiple users.

For more information, see the following sections:

Connecting to an Xgrid Controller

You use Xgrid Admin to connect to an Xgrid controller. The controller must be reachable on any network by the administrative computer running Xgrid Admin.

After Xgrid Admin is connected to the controller, you can view the status of its grid and manage its agents and jobs.

To connect to an Xgrid controller:

- 1 Open Xgrid Admin and do one of the following:
 - From the pop-up menu, choose the controller or enter its name and click Connect.
 - In the Controllers and Grids list, select the controller name and click Connect.
- 2 If necessary, select the correct authentication option, enter a password, and then click OK.

Disconnecting from an Xgrid Controller

You use Xgrid Admin to disconnect from an Xgrid controller in the Controllers and Grids list.

To disconnect an Xgrid controller:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select a controller.
- 3 Click Disconnect.

Adding an Xgrid Controller

You use Xgrid Admin to add an Xgrid controller to the Controllers and Grids list.

To add an Xgrid controller to the monitoring list:

- 1 Open Xgrid Admin.
- 2 Click Add Controller.
- 3 From the pop-up menu, choose a controller or enter its name and click Connect.
- 4 If necessary, select the correct authentication option, enter a password, and then click OK.

Removing an Xgrid Controller

You can easily remove an Xgrid controller from the Controllers and Grids list in Xgrid Admin.

To remove an Xgrid controller:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select a controller.
- 3 Click Remove Controller.

Managing Agents

Use Xgrid Admin to view, add, or delete agents. Xgrid Admin also uses status indicators to display the status of agents.

Although Server Admin provides a simple interface for enabling Xgrid services on one server or across a rack of Xserve systems, it doesn't provide a way to configure Xgrid on desktop computers running Mac OS X v10.3 or later.

If you are relying on volunteers to provide desktop agents, you can send instructions for enabling Xgrid from the Sharing pane of System Preferences.

If the volunteers are using Mac OS X v10.3, you must first download the Xgrid Agent for Mac OS X v10.3 and then use the Xgrid pane of System Preferences. You can download the Xgrid Agent for Mac OS X v10.3 from:

www.apple.com/server/macosx/xgrid.html

If you administer a group of computers and want the computers to participate in a grid using Xgrid, you can use the following methods:

- Apple Remote Desktop
- SSH
- NetBoot or NetInstall

Apple Remote Desktop

Apple Remote Desktop (ARD) v2.1 is a separate product available from Apple that integrates common administrative tasks across multiple computers (such as screen sharing, software installation, running UNIX scripts, and so on).

You can use ARD to remotely run System Preferences on each computer but it is usually simpler to change the preferences once and then push the new preferences file (`/Library/Preferences/com.apple.xgrid.agent.plist`) to all relevant nodes.

For more information, see the *Apple Remote Desktop Administration* guide at www.apple.com/server/documentation.

SSH

If you don't have ARD but you've set up SSH logins, you can do the same thing as ARD using the `scp` command-line tool (or `rsync`, if you've set that up). You can also use the `xgridctl` tool with the following command:

```
$ ssh root@remotehost xgridctl agent start
```

For more details, see the man pages for SSH, SCP, SFTP, or `rsync` in the Terminal application.

NetBoot or Network Install

For large networks, it often makes sense to use a common system image that is mounted or installed by each agent to configure the agents.

Although Xgrid isn't reason enough to use NetBoot, consider whether using Network Install would simplify your general administrator's tasks. If you use Netboot with Xgrid, all agents must have unique hostnames and must keep all files intact between reboots. For more information, see *System Imaging and Software Update Administration* at www.apple.com/server/documentation.

Viewing a List of Agents

You can see a list of agents for a controller in Xgrid Admin.

To see a list of agents for an Xgrid controller:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the grid.
- 3 Click Agents.
- 4 Select an agent in the list to see information about the CPU power and processors it uses.

The color bubble to the left of the name shows each agent's status. For details, see "Status Indicators in Xgrid Admin" on page 40.

Adding an Agent

You can add an agent to a controller in Xgrid Admin. You can add agents that are offline. The agents will be available to the controller when the computers are online or when the controller administrator makes the agents active.

To add an agent:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the controller.
- 3 Click Agents.
- 4 Click the Add (+) button below the list of agents.
- 5 Enter a name for the agent and click OK.

The agent is added to the list. The color bubble to the left of the name shows the agent's status. For details, see "Status Indicators in Xgrid Admin" on page 40.

Deleting an Agent

You can delete an agent for an Xgrid controller in Xgrid Admin.

To delete an agent:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the controller.
- 3 Click Agents.
- 4 Click the Delete (–) button below the list of agents.

Note: If you delete an agent that you know is on the local subnet and is configured to attach to that controller, wait a few moments and it will reappear in the list. If the agent doesn't reappear, use the Add (+) button and enter its name to retrieve it.

Managing Jobs

You use Xgrid Admin to manage jobs after they are submitted by a client.

You cannot move a job between grids.

Viewing a List of Jobs

You can see a list of jobs in Xgrid Admin.

To see a list of jobs:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the controller.
- 3 Click Jobs.
- 4 Select a job in the list to see details of that job.

Stopping a Job

You can stop a job in Xgrid Admin.

To stop a job:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the controller.
- 3 Click Jobs.
- 4 Select the job you want to stop.
- 5 Click the Stop button below the list of jobs.

Repeating or Restarting a Job

You can repeat a job or restart a stopped job in Xgrid Admin.

To repeat or restart a job:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the controller.
- 3 Click Jobs.
- 4 Select the job you want to repeat or restart.
- 5 Click the Start button below the list of jobs.

Deleting a Job

You can delete a job in Xgrid Admin.

To delete a job:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the controller.
- 3 Click Jobs.
- 4 Select the job you want to delete.
- 5 Click the Delete (-) button below the list of jobs.

Adding a Grid

You use Xgrid Admin to add a grid to an Xgrid controller in the Controllers and Grids list.

To add a grid:

- 1 Open Xgrid Admin.
- 2 Select the Xgrid controller you want to add the grid to.
- 3 Click the Add (+) button below the Controller and Grids list.
- 4 In the pop-up menu, enter a name for the new grid and click OK.

Deleting a Grid

You use Xgrid Admin to remove a grid from an Xgrid controller in the Controllers and Grids list.

To delete a grid:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the grid.
- 3 Click the Action pop-up menu below the Controller and Grids list and select Remove Grid.
- 4 Click OK.

Monitoring Grid Activity

You can quickly view the activity of a grid in Xgrid Admin. You can also view agents and job activity using Xgrid Admin. For more information, see “Viewing a List of Agents” on page 43 and “Viewing a List of Jobs” on page 44.

To monitor the activity of a grid:

- 1 Open Xgrid Admin.
- 2 In the Controllers and Grids list, select the Xgrid controller.
- 3 Click Overview to see the number of agents, the amount of processor power available and used, and the number of jobs running and pending.

Use this chapter to learn how to use Xgrid command-line tools and the Terminal application to submit jobs to a grid and to get information about jobs.

After you configure an Xgrid controller and add agents to a grid, you can use the Terminal application to send a job to the grid.

Structuring Jobs for Xgrid

Carefully planning and structuring a job can result in efficient use of the grid. For example, the best structure for a job that requires multiple searches of a large database may be to divide the database into multiple sections and provide a section to each agent in the grid.

About Job Styles

Different styles of jobs often require different handling. Similarly, the way a job is structured influences how efficiently the grid completes it.

Consider the following job styles:

- Everything in one single large job, with numerous small tasks.
- Everything divided into medium-sized jobs, where each job has roughly as many tasks as there are nodes in the grid. (This type of job is usually created by a meta job script, which divides the job into smaller chunks, each of which is a job in itself.)
- An entire workflow composed of several interrelated jobs.

Deciding how to structure a job can involve experimentation to discover the best way to complete it.

For example, you might create a simple, small version of a job in two styles, such as by planning all tasks in one job or by subdividing into multiple tiny jobs. Running both experimental jobs under similar conditions in the grid will give you a good idea of which job style is better suited to those conditions.

About Job Failure

Xgrid jobs can rely on message-passing interface (MPI) APIs. For jobs that rely on MPI, if a single task fails, the entire job fails and must be resubmitted. Therefore you should not use MPI-based jobs on grids with high task-failure rates.

Jobs that are more parallel in nature are generally unaffected by occasional task failures. Tasks are typically reassigned to other available agents to complete the job. Most jobs fall into this category.

Submitting a Job

You submit jobs to a grid using the command-line tool and Terminal. Example code is available on the Apple developer website (developer.apple.com) for alternative methods of submitting jobs. Also if you have Developer Tools installed you can view the examples located in `/Developer/Examples/Xgrid/`.

For more information about the syntax and options for the `xgrid` command-line tool, see the `xgrid` man pages.

Some developers and organizations offer specialized applications for submitting jobs to a grid. Or you can create such an application using Apple's developer tools for Xgrid.

When determining whether to use the `xgrid` command-line tool or another method for submitting jobs, consider these points:

- If the job is simple, use the command-line tool.
- If you use a shell script, use the command-line tool.
- If you want to use Xgrid as part of an application with a graphical user interface (GUI), use the Xgrid API to create the GUI or incorporate it in an existing application. For more information about the API, see the *Xgrid Reference* at: developer.apple.com/documentation

Examples of Xgrid Job Submission and Results Retrieval

The following Terminal commands are examples of jobs a client can submit to the controller.

```
$ xgrid -h <controller> -p <password> -job submit /bin/echo "Hello, World!"
```

This job runs `/bin/echo` on the controller and agent systems with the "Hello, World!" parameter.

```
$ xgrid -h <controller> -p <password> -job results -id <id>
```

This command shows the results of the job with the id indicated.

For an executable shell script marked `hello.sh`:

```
#!/bin/sh
/bin/echo "Hello, World!"
```


The following command copies the shell script `hello.sh` to the Xgrid controller and agent systems and runs the script. `/bin/echo` must be installed on the agent system. The `hello.sh` script must have its executable bit set before it can execute.

```
xgrid -h <controller> -p <password> -job submit hello.sh
```

Viewing Job Status

You can monitor jobs in Xgrid Admin (for details, see “Managing Jobs” on page 44) or with the command-line tool.

The following commands in Terminal provide job status:

```
$ xgrid -h <controller> -p <password> -job list
$ xgrid -h <controller> -p <password> -job attributes -id <job-id>
```

Retrieving Job Results

You can retrieve job results using the command-line tool.

The following commands in Terminal retrieve job results.

```
$ xgrid -h <controller> -p <password> -job results
$ xgrid -h <controller> -p <password> -job results id <job-id>
```


Use this chapter to help solve common problems you might encounter and questions you might have while working with Xgrid service.

This section contains answers to common problems and questions.

If Your Agents Can't Connect to the Xgrid Controller

If an agent is a server, make sure the agent service is enabled and the Xgrid service is started. The Xgrid controller is the only component of Xgrid that has an open port (port 4111) and requires a firewall opening.

This means the Xgrid controller is the only component that advertises on or responds to queries over Bonjour. When enabling the controller, make sure firewall port 4111 is open on your computer's firewall (enabled in the Sharing Pane of System Preferences) or your corporate firewall (if accepting agents or clients outside your organization).

Agents and clients access the controller through a Bonjour lookup or an explicit hostname/IP address, then they initiate a connection to the controller over a user port, avoiding the need to perform privileged operation or opening the firewall.

If You Use Xgrid over SSH

The simplest way to secure Xgrid using SSH is to create a tunnel from the client or the agent to the controller:

```
$ ssh user@controller.hostname.com -L 4111:controller.hostname.com:4111
```

Then, have the agent or client connect to localhost instead of the controller. By doing this, SSH tunnels to the remote connection. You can use other ports on the local machine and even tunnel through an intermediary host.

To run an Xgrid agent over an SSH tunnel as a particular user:

Using Terminal, enter the following:

```
$ ssh -R 20000:192.168.1.100:4111 user@192.168.1.102 /usr/libexec/xgrid/  
GridAgent -ServiceName localhost:20000 -RequireControllerPassword NO -  
UsesRendezvous NO -OnlyWhenIdle NO -BindToFirstAvailable NO
```

20000 is the port to tunnel through the ssh connection, *192.168.1.100:4111* is the address and port number of the controller, *user* is the name of the user to connect, and *192.168.1.102* is the address of the remote computer to run the agent.

If You Run Tasks on Multi-CPU Machines

By default, each Xgrid agent (one per machine) accepts as many tasks as there are CPUs on that host, as reported by `$ sysctl hw.ncpu`.

Agents assume that tasks are single-threaded, so they will run two tasks to make best use of a dual-CPU system. To run multithreaded tasks that take up both CPUs, edit the agent configuration file `/Library/Preferences/com.apple.xgrid.agent.plist`.

To make it always only accept a single task, change the `MaximumTaskCount` line to:

```
MaximumTaskCount=1
```

Note: This must be done explicitly for each agent, and is permanent until reversed. You can't specify this kind of constraint as part of a job submission.

If You Submit a Large Number of Jobs

GridStuffer is a third-party Cocoa application created by Charles Parnot of Stanford to manage multitask jobs. It provides a friendly GUI for many common Xgrid tasks.

GridStuffer is available at:

<http://cmgm.stanford.edu/~cparnot/xgrid-stanford/html/goodies/GridStuffer-info.html>

A companion command-line tool, `xgridstatus`, provides an easy way to retrieve information about your grid and jobs. `Xgridstatus` is available at:

<http://cmgm.stanford.edu/~cparnot/xgrid-stanford/html/goodies/xgridstatus-info.html>

If You Want to Use Xgrid on Other Platforms

Third-party agents are available that run Xgrid jobs on non-Mac platforms. You are responsible for ensuring that your tasks contain and call appropriate platform-specific code.

There is no intrinsic support for heterogeneous execution, although there is nothing that relies on Mac-specific technology.

The primary technical requirement is a sufficiently functional BEEP protocol stack. Several open source implementations are available, of varying quality.

Two cross-platform Xgrid agents are available:

- Curtis Campbell's java agent, at:
<http://sourceforge.net/projects/xgridagent-java/>
- Daniel Cote's Linux/UNIX agent (not yet updated for Mac OS X v10.4), at:
<http://www.novajo.ca/xgridagent/>

If the Xgrid Controller Must Be Restarted

When the Xgrid controller is restarted, by Server Admin, `xgridctl` tool, a power-outage, or a kernel panic, the following occurs:

- Clients and agents are disconnected.
- Tasks running when the controller restarted are stopped.
- Partial data from killed tasks is discarded.
- data from finished tasks is saved and can be retrieved as usual.
- Queued jobs and tasks are saved and run as usual.
- Tasks are started/restarted as agents reconnect and become available.

If Xgrid Has Crashed

The Xgrid controller and agent should restart automatically if they crash. CrashReporter logs can be found in `/Library/Logs/CrashReporter`. Xgrid logs notices, warnings, and errors to the console as well as to log files in `/Library/Logs/Xgrid`

If You Are Trying to Submit Jobs over 2 GB

The Xgrid controller is a 32-bit process and keeps most job input and output data in memory. This means that the controller can crash if your jobs require a large amount of input or produce a large amount of output. This limitation might change in the future.

We recommend using a shared filesystem (such as Xsan or NFS) if you need to share large amounts of data between distributed processes.

If You Want to Enable Kerberos/SSO for Xgrid

For Xgrid to use SSO, you need the following:

- The agent *must* have the host's user principal in the system keytab.
- The Kerberos database on the KDC *must* contain the agent's principal.
- The controller's realm *must* be the default realm on the agent computer.

The agent's principal is created in the KDC and is put in the agent's keytab if the agent computer is bound to the OD master using `_AUTHENTICATED BINDING_` with Directory access. Otherwise, you must use `kadmin` to create the principal in the KDC and export it to the keytab.

For example, the computer hosting the agent must have the host's user principal in the system keytab, as shown here:

```
$ hostname:~ user
$ sudo klist -k
$ Password:
$ Keytab name: FILE:/etc/krb5.keytab
KVNO Principal
-----
  1 hostname.apple.com@XGRIDTEST.APPLE.COM
  1 hostname.apple.com@XGRIDTEST.APPLE.COM
  1 hostname.apple.com@XGRIDTEST.APPLE.COM
```

The Kerberos database on the KDC must contain the agent's principal, as in the following:

```
$ sudo kadmin.local -q "get_principal hostname.apple.com"
Authenticating as principal root/admin@XGRIDTEST.APPLE.COM with password.
Principal: hostname.apple.com@XGRIDTEST.APPLE.COM
Expiration date: [never]
Last password change: Tue Apr 12 17:46:41 PDT 2005
Password expiration date: [none]
Maximum ticket life: 0 days 10:00:00
Maximum renewable life: 7 days 00:00:00
Last modified: Tue Apr 12 17:46:41 PDT 2005 (root/admin@XGRIDTEST.APPLE.COM)
Last successful authentication: [never]
Last failed authentication: [never]
Failed password attempts: 0
Number of keys: 4
Key: vno 1, Triple DES cbc mode with HMAC/sha1, no salt
Key: vno 1, ArcFour with HMAC/md5, no salt
Key: vno 1, DES cbc mode with CRC-32, no salt
Key: vno 1, DES cbc mode with CRC-32, Version 4
Attributes: REQUIRES_PRE_AUTH
Policy: [none]
```

The controller's realm must be the default realm on the agent computer, as shown:

```
$ cat /Library/Preferences/edu.mit.Kerberos
# WARNING This file is automatically created, if you wish to make changes
# delete the next two lines
# autogenerated from : /LDAPv3/xgridtest.apple.com
# generation_id : 1637891359
[libdefaults]
    default_realm = XGRIDTEST.APPLE.COM
[realms]
    XGRIDTEST.APPLE.COM = {
        kdc = xgridtest.apple.com
        admin_server = xgridtest.apple.com
    }
[domain_realm]
    apple.com = XGRIDTEST.APPLE.COM
    .apple.com = XGRIDTEST.APPLE.COM
```

For More Information

If you're an experienced server administrator or even a novice server administrator working with Xgrid, you can review the Xgrid FAQ site. The FAQ site will provide you with access to news, posted questions and threads, and the ability to post your own Xgrid questions.

The site is at http://lists.apple.com/faq/pub/xgrid_users/.

For more information about advanced configuration options, see the `xgridctl` man page.

Part II: Configuring High Performance Computing



Use the chapters in this part of the guide to learn about high performance computing and the applications and tools available for administering it.

Chapter 6	Introducing High Performance Computing
Chapter 7	Reviewing the Cluster Setup Process
Chapter 8	Identifying Prerequisites and System Requirements
Chapter 9	Preparing the Cluster for Configuration
Chapter 10	Setting Up the Cluster Controller
Chapter 11	Setting Up Compute Nodes
Chapter 12	Testing Your Cluster

Introducing High Performance Computing

6

Use this chapter to learn about high performance computing (HPC) and how it's supported by Apple technology.

With high performance computing, you can speed the processing of complex computations by using Xserve computers with the Xgrid service.

Understanding HPC

HPC refers to the use of high-end computer systems to solve computationally intensive problems. HPC includes large supercomputers, symmetric multiprocessing (SMP) systems, cluster computers, and other hardware and software architectures.

In recent years, developers have made it feasible for standard off-the-shelf computer systems to achieve supercomputer-scale performance by clustering them in efficient ways.

Apple and HPC

Apple's hardware and software facilitate HPC in unique and meaningful ways. Although many hardware and software architectures can be used for cluster computing, Mac OS X Server v10.5 and Xserve have specific features that enhance the performance and manageability of cluster installations.

The integration of Xserve with Mac OS X Server provides unparalleled ease of use, performance, and manageability. Because Apple makes the hardware and the software, the benefits of tight integration are immediately evident in the quality of the user experience with a Macintosh-based cluster.

Mac OS X Server

Mac OS X Server v10.5 is Apple's award-winning UNIX server operating system. Mac OS X Server can compile and run UNIX 03-complaint code, and runs 64-bit applications alongside 32-bit applications at native performance.

The Mach kernel provides preemptive multitasking for outstanding performance, protected system memory for stability, and modern SMP locking for efficient use of multi processor and multi core systems.


Mac OS X Server also includes highly optimized math libraries that enable software developers to take maximum advantage of the G5 or Intel-based processor without the use of difficult programming techniques or expensive development tools.

Mac OS X Server also includes Xgrid, an integrated distributed resource manager for both grids and clusters.

Xserve Clusters

Using a combination of Xserve systems, you can build clusters that aggregate the power of these systems to provide HPC solutions at comparatively low cost.

An Xserve cluster consists of at least 2 nodes: a cluster controller and one or more compute nodes, as shown in the following illustration:



Xserve 64-Bit Architecture

The 64-bit architecture of Xserve systems is ideal for HPC applications. It provides 64-bit math precision, higher data throughput, and very large memory space.

Memory Space

The 64-bit architecture provides four billion times the memory space available in a 32-bit architecture, which puts the theoretical address space available to Mac OS X Server applications at 16 exabytes. Xserve G5 systems support 8 GB of memory. Xserve Intel systems support 32 GB of memory.

Libraries

Mac OS X Server provides the following highly optimized libraries for developing HPC applications. In addition to standard libraries like libSystem, numerical libraries like BLAS, LAPACK, and others provide industry-standard routines that have been hand-tuned for the G5 or Intel processor. Developers can make efficient use of the system architecture without writing computer code or vector code.

Library	Description
libSystem	A collection of core system libraries
libMathCommon	A common math functions library
vDSP	A library that provides mathematical functions for applications that operate on real and complex data types
BLAS	A library of basic linear algebra subprograms, which are a standard set of building blocks for vector and matrix operations
LAPACK	The linear algebra package, which is a standard library for solving simultaneous linear equations
vForce	A library of highly-optimized single- and double-precision mathematical intrinsic functions
vBasicOps	A collection of basic operations that complement the vector processor's basic operations up to 128 bits
vBigNum	A library of optimized arithmetic operations for 256-, 512-, and 1024-bit operands

Easy Porting of UNIX Applications

Mac OS X Server is now an Open Brand UNIX 03 Registered Product, conforming to the SUSv3 and POSIX 1003.1 specifications for the C API, shell utilities, and threads. It can compile and run all your existing UNIX 03-compliant code.

Support of Loosely Coupled Computations

You can use Xserve clusters to perform most types of loosely coupled or *embarrassingly parallel* computations. Embarrassingly parallel computations consist of somewhat independent computational tasks that can be run in parallel on many different processors to achieve faster results.

Here are examples of loosely coupled computations that you can accelerate using the setup described in this guide:

- **Image rendering.** Different rendering tasks, such as ray tracing, reflection mapping, and radiosity, can be accelerated by parallel processing.
- **Bioinformatics.** The throughput of bioinformatics applications like BLAST and HMMER can be greatly enhanced by running them on a cluster.

Note: The Apple Workgroup Cluster is a preconfigured cluster solution that has everything you need to get up and running quickly. It includes qualified, integrated hardware components and easy-to-use management tools. You can add cluster-aware commercial applications, such as iNquiry or gridMathematica, or develop your own custom applications using Xcode. For more information, see <http://www.apple.com/science/solutions/workgroupcluster.html>.

- **Cryptography.** Brute-force key search is a classic example of a cryptography application that can be greatly accelerated when run on a computer cluster.
- **Data mining.** High performance computing is essential in data mining because of the amount of data that is analyzed.

Note: This guide assumes that the cluster nodes communicate over gigabit Ethernet. Although the network latency of Gigabit Ethernet is low enough for most loosely coupled computations, those that require lower latency may benefit from another interconnect technology.

Reviewing the Cluster Setup Process

7

Use this chapter to learn about the process of setting up a high performance cluster.

You will use multiple server tools to configure services, a cluster controller, compute nodes, and users when setting up a high performance cluster.

The following chapters provide a step-by-step process to assemble and configure a computational cluster. The resulting cluster will consist of a controller and a number of compute nodes. The compute nodes will be connected to the controller via a private (isolated) Ethernet network switch. The controller will be connected to both the private Ethernet network and a public network, potentially the Internet. The controller will also provide a shared file system to compute nodes.

The controller will provide a number of services to the compute nodes:

- A Firewall will isolate the controller and compute nodes from the public network, protecting against unwanted access. Access to the private network from outside the firewall will require remote users to use SSH for command-line access or VPN to use or manage cluster resources with graphical applications or administrative tools such as Apple Remote Desktop.
- Network services such as DHCP, DNS, and NAT will allow the compute nodes to communicate with each other and external networks.
- Open Directory will contain user account information, including usernames and passwords, and make these accounts available to compute nodes. Using Kerberos with Open Directory provides single sign-on capability, reducing the number of times a user will need to enter passwords to access cluster resources.
- Open Directory will also publish network file system (NFS) share points, providing automatic file sharing between compute nodes and controller. A shared network home directory, containing home folders for each cluster user, will be mounted on each compute node.
- The controller will host the Xgrid controller service.

Cluster Setup Overview

Here is a summary of what you'll be doing to set up and test an HPC cluster.

Step 1: Before you begin

Before setting up your cluster, understand the expectations and requirements that you must fulfill. See Chapter 8, "Identifying Prerequisites and System Requirements."

Step 2: Prepare the cluster for configuration

Prepare your cluster nodes for configuration by setting up the hardware and connecting your nodes to a network. See Chapter 9, "Preparing the Cluster for Configuration."

Step 3: Enable, configure, and start services

After your cluster is assembled and ready, start by setting up and configuring the cluster controller. Use Server Assistant to set up the server software on the cluster controller. See Chapter 10, "Setting Up the Cluster Controller."

Use Server Admin to configure and start the following services:

- DNS service. See "Configuring DNS Service" on page 84.
- Open Directory service. See "Configuring Open Directory Service" on page 86.
- DHCP service. See "Configuring DHCP Service" on page 87.
- Firewall service. See "Configuring Firewall Settings on the Cluster Controller" on page 88.
- NAT service. See "Configuring NAT Settings on the Cluster Controller" on page 90.
- NFS service. See "Configuring NFS" on page 90.
- VPN service. See "Configuring VPN Service" on page 90.
- Xgrid service. See "Configuring Xgrid Service" on page 91.

Step 4: (Optional) Prepare the data drive

Use Disk Utility to configure the data drive. See "Preparing the Data Drive as a Mirrored RAID set" on page 92.

Step 5: Create an automounted network share

Use Server Admin to create an automounted network share. See "Creating a Home Directory Automount Share Point" on page 93.

Step 6: Create network user accounts

Use Workgroup Manager to create network user accounts for cluster users. See "Creating User Accounts" on page 94.

Step 7: Create an Auto Server Setup record for the compute nodes

Use Server Assistant to save configuration settings to a file or Open Directory record. This allows cluster nodes to automatically configure themselves when they start up for the first time.

See “Creating an Auto Server Setup Record for Compute Nodes” on page 95 and “Verifying LDAP Record Creation” on page 98.

Step 8: Set up compute nodes

Start compute nodes to begin the Auto Server Setup process. They’ll automatically configure themselves and then restart. See “Setting Up Compute Nodes” on page 98.

Step 9: Finish compute node configuration

Use Server Admin to name the compute nodes, join them to the Kerberos realm, and configure their Xgrid agent software.

Step 10: Test your cluster setup

After configuring the controller and compute nodes, test your cluster with Xgrid Admin and a sample Xgrid application. See Chapter 12, “Testing Your Cluster.”

Identifying Prerequisites and System Requirements

8

Before setting up your cluster, read the prerequisites and requirements in this chapter and familiarize yourself with the setup process.

To make sure that your cluster is successfully set up, read this chapter to familiarize yourself with the expectations and requirements you must meet before starting the setup procedure. Then read the last section, which provides an overview of the cluster setup process.

Prerequisites

This guide assumes you have the expertise needed to set up and manage the cluster, perform the initial configuration of the cluster nodes, and carry out the types of computations you can perform on the cluster.

Expertise

To set up and deploy clusters, you should have a good understanding of how Mac OS X Server works and you should have a fundamental understanding of UNIX, Xgrid, and TCP/IP networking.

Xserve Configuration

This guide assumes that you'll be using new, out-of-the-box Xserve systems running Mac OS X Server v10.5 or later. If not, you must install a clean version of Mac OS X Server v10.5 or later on your systems.

System Requirements

Take time to define the requirements needed to make sure the cluster setup is successful. System requirements are categorized as infrastructure, software, and private network requirements.

Infrastructure Requirements

This section describes the most important hardware infrastructure requirements. Consult with your system administrator about other requirements.

For example, you might need one or more uninterruptible power supplies (UPSs) to provide backup power to key cluster components. Another requirement might be a physical security system to protect the cluster from unauthorized access to sensitive information.

Infrastructure requirements are divided into the following subcategories:

- “General Hardware Requirements” on page 68
- “Power Requirements” on page 68
- “Cooling Requirements” on page 69
- “Weight Requirements” on page 70
- “Space Requirements” on page 70
- “Network Access Requirements” on page 71

General Hardware Requirements

To set up a cluster, you should have the necessary hardware infrastructure in place. This includes:

- Racks
- Electrical power
- Cooling system
- Network access points and switches

Power Requirements

When setting up the physical infrastructure for your cluster, consider the following power consumption figures:

- **Rated power consumption.** This figure represents the *maximum* power consumption of a given system’s power supply.
- **Typical power consumption.** This figure represents the *typical* power consumption of a server under normal operating conditions.

Note: This section focuses only on the rated power consumption figure because it guarantees that your circuit won’t be overloaded at any time—unlike the typical power consumption figure, which doesn’t protect your circuit from abnormal surges in power consumption.

To obtain power consumption figures for cluster nodes, see the following articles on the AppleCare Service & Support website:

- Article 86694, "Xserve G5: Power consumption and thermal output (BTU) information," at www.info.apple.com/kbnum/n86694
- Article 75383, "Xserve: Power Consumption and Thermal Output (BTU) Information," at www.info.apple.com/kbnum/n75383
- Article 86251, "Xserve (Slot Load): Power Consumption and Thermal Output (BTU) Information," at www.info.apple.com/kbnum/n86251
- Article 304887, "Xserve (Late 2006): Power consumption and thermal output (BTU) information," at www.info.apple.com/kbnum/n304887

Although the rated current load covers your cluster nodes, you must also consider the power consumption of other devices connected to your circuit.

For large clusters, speak with an Apple Systems Engineer to determine the correct power infrastructure. For information about Apple consulting services and service and support plans, see the Apple Server Service and Support website at <http://www.apple.com/server/support>.

WARNING: The formulas in this section help you estimate your power requirements. These estimates may not be high enough, depending on your configuration. For example, if your cluster uses one or more Xserve RAID systems, or other third-party hardware, you must include their power consumption requirements.

Cooling Requirements

It's very important to keep your Xserve computers running at normal operating temperatures (see www.apple.com/xserve/specs.html). If your servers overheat they will shut down and any work being done will be lost. You can also damage or shorten the life span of your servers by running them at high temperatures.

To obtain thermal output figures for cluster nodes, see the following articles on the AppleCare Service & Support website:

- Article 86694, "Xserve G5: Power consumption and thermal output (BTU) information," at www.info.apple.com/kbnum/n86694
- Article 75383, "Xserve: Power Consumption and Thermal Output (BTU) Information," at www.info.apple.com/kbnum/n75383
- Article 86251, "Xserve (Slot Load): Power Consumption and Thermal Output (BTU) Information," at www.info.apple.com/kbnum/n86251
- Article 304887, "Xserve (Late 2006): Power consumption and thermal output (BTU) information," at www.info.apple.com/kbnum/n304887

Consider the thermal output of other devices, such as the management computer, Xserve RAID systems, monitors, and other heat-generating devices used in the same room.

As always, consult with your system administrator to determine the necessary level of cooling that your cluster and its associated hardware require for safe and effective operation.

Weight Requirements

For Xserve and cluster node weight information, see the Apple Xserve website at www.apple.com/xserve.

Also include the weight of the rack if you're bringing in a dedicated rack, and the weight of other devices used by the cluster.

If you mount cluster nodes in a rack with casters, set up the rack where you'll keep the cluster and then mount the systems. A heavy rack is difficult to move, particularly across carpet. In addition, vibrations caused by moving your cluster long distances when racked might damage your hardware.

After determining weight requirements, consult with your facilities personnel to make sure the room where the cluster will be installed meets the weight requirements.


Space Requirements

You should have enough space to house the cluster and enable easy access to it to perform routine maintenance tasks. Also, locate the cluster where it doesn't affect and isn't affected by other hardware in your server room.

Consider the following when choosing a location for your cluster:

- Don't place the cluster next to an air vent, air intake, or heat source.
- Don't place the cluster directly under a sprinkler head.
- Don't obstruct doors (especially emergency exit doors) with your cluster.
- Leave enough room in front of, beside, and especially behind your cluster.
- Make sure air can flow around the cluster. The room might be very well cooled, but if air can't easily flow around the cluster, your computers can still overheat.

If you're housing your cluster in a computer room, make sure you have at least 18 inches of clearance in front and behind your systems. If you're housing it in an office or other unmanaged space, make sure your cluster has at least 18 inches of clearance on all sides of the rack, as shown in the following illustration:



You should have enough space to open the rack's door, slide out systems, and perform other routine maintenance tasks.


Network Access Requirements

Your cluster requires access to two networks:

- **Private network.** This is a high performance Gigabit Ethernet network. You'll need at least a 1-Gigabit switch.
- **Public network.** This network connects the cluster controller to the client computers that submit jobs to your cluster.

This guide uses a number of 10.0.2.x addresses as examples for your public network connections. Do not use these example addresses when configuring your cluster. When you see a 10.0.2.x address, substitute the IP address appropriate for your organization's network.

The following illustration shows a configuration of a cluster connected through a switch creating a private network. The illustration also shows the headnode connected to the public and private network.



Software Requirements

You need:

- A site-licensed copy of Mac OS X Server v10.5 or later.
- One or more copies of Apple Remote Desktop v3 or later (recommended).
- The latest version of Server Tools.

Volume-Licensed Serial Number

To run multiple copies of Mac OS X Server, you should obtain a volume-licensed serial number. If you haven't obtained a volume-license serial number, contact your local Apple sales representative.

Note: The format of the server serial number is xsvr-999-999-x-xxx-xxx-xxx-xxx-xxx-x, where x is a letter and 9 is a digit. The first element (xsvr) and the fourth (x) must be lowercase.

Apple Remote Desktop

Configuration and administration of your cluster will be greatly enhanced with Apple Remote Desktop v3 or later. You can use Apple Remote Desktop to configure, monitor, and control your cluster, as well as rapidly install software.

Server Tools

If you are using a management computer, you must install Server Tools on your management computer. The Server Tools suite includes:

- Server Assistant
- Server Admin
- Server Monitor
- Xgrid Admin

You use these tools to remotely manage the cluster. Install these tools using the Server Admin Tools CD, which is included with Xserve and Mac OS X Server.

Private Network Requirements

The compute nodes will be connected through a private Ethernet network, separate from your organization's primary (public) network. The cluster controller will be connected to the private and public networks and will act as a gateway, allowing users connected to the public network (or the Internet) to use the cluster's resources, and allowing the compute nodes to use resources outside the private network.

Private network requirements include the following:

- A range of IP addresses should be reserved for the private network. A number of non-routable IP address ranges are available for use with private networks. These addresses cannot be used with the Internet without Network Address Translation (NAT), which will be provided by the cluster controller.

- Addresses in ranges such as 192.168.x.x, 10.0.x.x, and 172.16.x.x are commonly used for private networks. Because the first two are used more commonly with NAT devices used in the home, and because your users may want to connect to your cluster from behind one of these devices, it is best to choose a range less likely to exist on your user's networks. This guide uses the range 172.16.1.1 - 172.16.1.254 (subnet mask 255.255.255.0). You can use this range for your cluster, or use a different one if you prefer.
- You need a Domain Name System (DNS) server that will be used to assign names to network addresses so you don't need to remember IP addresses. Your private network can use a DNS domain name that is not in use on (and is not valid with) the Internet. This guide uses the .cluster domain. You can use this domain with your cluster as well.

WARNING: Where you see the DNS domain .example.com, you should substitute the DNS domain used for your organization's public network.

Static IP Address and Hostname Requirements

Your cluster requires a single static IP address and a matching fully qualified and reverse resolvable DNS entry for the cluster controller.

By using a static IP address rather than a dynamic one you can maintain a consistent address that clients can always use.

Note: Initiate the process of requesting an IP address and a hostname as early as you can before setting up the cluster, to account for the lead time typically required.

Preparing the Cluster for Configuration

9

Use this chapter to mount the systems on the rack, connect the systems to a power source and the private network, and configure the optional management computer.

To prepare the cluster nodes for configuration, you mount them in racks and connect them to the power source and private network. You also set up the management computer by installing Apple Remote Desktop and Server Tools.

Preparing the Cluster Nodes for Software Configuration

After you prepare the physical infrastructure for hosting the cluster, the next step is to mount the cluster nodes and prepare them for software configuration.

To prepare the cluster for configuration:

- 1 Unpack the computers and mount them in the rack.

For more information, use the instructions provided with your hardware.

Note: If you're using existing Xserve computers, you must perform a clean installation of Mac OS X Server v10.5 or later to restore the systems to default settings.


- 2 Record each computer's serial number and keep the information in a safe place.

When recording the serial numbers, do it in a way that makes it easy for you to tell which serial number belongs to each computer. For example, use a table to map a system's serial number to the name on a label on the system's front panel.

Serial Number	Name
<i>serial_number_0</i>	Cluster controller
<i>serial_number_1</i>	Compute node 1
<i>serial_number_2</i>	Compute node 2
...	...


You can find the serial number of an Xserve computer in four places:

- The unit's back panel:



- The unit's interior

If you look for the serial number on the unit's interior, don't confuse the serial number for the server with the serial number for the optical drive—these are different numbers. The Xserve computer's serial number is denoted by "Serial#" (not "S/N") followed by 11 characters.




- The large pull-out plastic tab on Xserve computers with Intel processors
- The cardboard shipping box
You can use a barcode scanner on the box label to get the serial number.

3 Use the following guidelines, connect the cluster computers to a power source:

- **Power cables.** Use the long power cables with a horizontal power distribution unit (PDU) and the short cables with a vertical PDU. When using the long cables, connect the servers so you can tell which cable belongs to which node. Consider labeling cables to make it easier to map a cable to a node.
- **Connection to the uninterruptible power supply (UPS).** Connect the cluster controller, storage devices used by the cluster, and the private network switch to a UPS unit to protect against data loss in case of a power outage. If your UPS is connected to the controller through USB, you can use the UPS configuration settings in System Preferences.

Note: If you are using a UPS, the UPS low power shutdown script is available for additional advanced power options. This script is located at `/usr/libexec/upsshutdown`.

- **UPS connection to wall outlet.** Make sure the electrical outlets support the UPS plug shape.
 - **Power cord retainer clips.** To prevent power cables from slipping out, use the power cord retainer clips that come with your Xserve systems.
 - **Air flow.** Don't permit a mass of power cables to obstruct air flow.
- 4 Connect the two Ethernet ports (shown in the illustration below) by connecting port 1 on the cluster controller to the public network and port 2 to the private network.



- 5 Connect Ethernet port 1 on the remaining nodes in the cluster to the private network, in order.

Use the last port on the switch for the cluster controller, the first port for the first compute node, the second port for the second compute node, and so on.

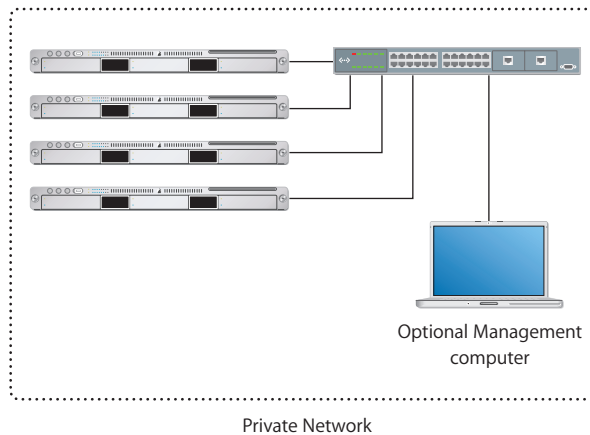
Connecting the Ethernet cables to the switch in order helps you identify which cluster node a cable belongs to.

(Optional) Setting Up the Management Computer

You can use the management computer to remotely set up, configure, and administer your cluster.

To set up the management computer:

- 1 Connect the management computer to the private network (as shown) using the second-to-last switch port.



- 2 Start the management computer.
- 3 Disable AirPort and any network connection other than the one you'll be using to connect to your private network.
- 4 If they aren't installed, install the latest version of the Mac OS X Server tools and applications from the *Mac OS X Server Administration Tools* CD, which is included with the Mac OS X Server installation kit.

The Mac OS X Server tools and applications are installed into `/Applications/Server/`.

- 5 Configure the management computer's network address.

If your cluster controller is not connected to a keyboard, video display, and mouse, or if you prefer to set up the cluster from a management computer, you will connect the management computer to the private network and disable all other network connections.

Until the controller is assigned an IP address on the private network, configure your management computer to use DHCP. After the controller is assigned an IP address, you should configure your management computer to use a static address in the range reserved for your private network, but outside the range reserved for compute nodes.

If you are adopting the IP address range that is used in this guide (172.16.1.1 - 172.16.1.199 for compute nodes, 172.16.1.254 for the controllers), you can configure your management computer to use 172.16.1.253.

After you connect to the private network, the server administration tools mentioned in this guide (Server Assistant, Server Admin, Workgroup Manager, and Xgrid Admin) can be installed and used on your management computer, connecting via IP address to the cluster controller (and later the compute nodes).

You can also use Apple Remote Desktop, or the screen sharing feature included with Mac OS X v10.5, to control the nodes via the network, using the server administration tools directly on the remote nodes.

Use this chapter to set up server software on the cluster controller and configure the services running on it.

You use Server Assistant, Server Admin, and Apple Remote Desktop (optional) to set up and configure the cluster controller.

Setting Up Server Software on the Cluster Controller

To set up the cluster controller, use Server Assistant (located in /Applications/Server/).

To set up the cluster controller:

- 1 Start the cluster controller.

The cluster controller should have two Ethernet cables, with Ethernet port 1 connected to the public network switch and Ethernet port 2 connected to the private network switch. Only the cluster controller should be running on the private network.

If you are using a management computer, use Server Assistant to connect to the controller. For more information about using Server Assistant remotely, see *Server Administration*.

If you are using the Apple Remote Desktop to manage the controller, connect to the controller and initiate a screen control session. For more information, see the *Apple Remote Desktop Guide*.

- 2 In the Welcome screen, click Continue.
- 3 In the Server Configuration screen:
 - a Select Advanced.
 - b Click Continue.
- 4 In the Keyboard screen:
 - a Select the keyboard layout for the server.
 - b Click Continue.

- 5 In the Serial Number screen:
 - a Enter a volume license Mac OS X Server serial number.
 - b Click Continue.
- 6 In the Registration Information screen, fill out the form or press Command-Q and click Skip.
- 7 In the Administrator Account screen:
 - a Create the user account you'll use to administer the cluster controller (for example, Administrator).
 - b Click Continue.
- 8 In the Network Address screen:
 - a Choose "No, configure network settings manually."
 - b Click Continue.
- 9 In the Network Interfaces screen:
 - a Enable TCP/IP only for Ethernet 1 and Ethernet 2 by selecting the checkboxes for both Ethernet 1 and Ethernet 2.
 - b Click Continue.
- 10 In the TCP/IP Connection screen for the Ethernet 1 port:
 - a From the Configure pop-up menu, choose Manually.
 - b In the IP Address field, enter the public IP address of the cluster controller (for example, 10.0.2.199).
 - c In the Subnet Mask field, enter the public subnet mask of the cluster controller (for example, 255.255.255.0).
 - d In the Router field, enter the IP address of the router for the public network (for example, 10.0.2.1).
 - e Leave the DNS Servers field blank.
 - f Leave the Search Domains field blank.
 - g Click Configure IPv6.
 - h From the Configure IPv6 pop-up menu, choose Off.
 - i Click OK, then click Continue.
- 11 In the TCP/IP Connection screen for the Ethernet 2 port:
 - a From the Configure pop-up menu, choose Manually.
 - b In the IP Address field, enter the private IP address of the cluster controller (for example, 172.16.1.254).
 - c In the Subnet Mask field, enter the private subnet mask of the cluster controller (for example, 255.255.255.0).
 - d In the Router field, enter the private IP address of the cluster controller (for example, 172.16.1.254).

- e Leave the DNS Servers field blank.
 - f Leave the Search Domains field blank.
 - g Click Configure IPv6.
 - h From the Configure IPv6 pop-up menu, choose Off.
 - i Click OK, then click Continue.
- 12 In the Network Names screen:
- a Enter the primary DNS name and computer name.
The cluster controller has a public and a private DNS name. Use the controller's private names. For example, use controller.cluster for the primary DNS name and controller for the computer name.

A warning may appear saying the server's address resolves to another name. Click OK.
 - b Verify that the Enable Remote Management checkbox is selected.
 - c Click Continue.
- 13 In the Time Zone screen:
- a In the Closest City pop-up menu, choose your time zone.
 - b Click Continue.
- 14 In the Directory Usage screen:
- a From the "Set directory usage to" pop-up menu, choose Standalone Server.
 - b Click Continue.
- 15 In the Confirm Settings screen:
- a Review the settings.
 - b Click Apply.
 - c Wait for your settings to be applied.
- 16 Click Start Now, wait until Server Admin launches, and then (if prompted) enter the administrator user name and password.
- 17 When prompted, click Start Now; then when Server Admin launches, connect using the administrator user name and password.
- 18 Select the checkboxes to enable the following services: DHCP, DNS, Firewall, NAT, NFS, Open Directory, VPN, and Xgrid.
- 19 Click Save.
- 20 To reveal the enabled services, expand the triangle next to the controller in the Servers list.

Configuring DNS Service

Use Server Admin on the cluster controller to create a local DNS zone and add records to map cluster nodes to their corresponding IP addresses.

To configure DNS service:

- 1 Open Server Admin if it is not already open.
- 2 If necessary, click the triangle to the left of the controller to view a list of services.
- 3 Click DNS in the expanded Servers list.
- 4 Click Settings.
- 5 Click the Add (+) button below the “Forwarder IP Addresses” list, then enter the network address of your public DNS server (for example, 10.0.2.201).
- 6 Click Save.
- 7 Click Zones.
- 8 Click the Add Zone button, then select “Add Primary Zone (Master).”
A default zone named example.com is created.
- 9 Select the default example.com zone.
- 10 Change the primary zone name to your private DNS domain.
The primary zone name must end with a period (for example, “cluster.”).
- 11 Set Admin Email to the mail address of the person who should be notified of DNS errors (for example, administrator@example.com).
- 12 Double-click the first entry in the Nameservers list and change it to the private DNS hostname of the cluster controller (for example, controller).
- 13 Click Save.
- 14 Select the cluster DNS zone.
- 15 Click the triangle to the left of the cluster DNS zone.
- 16 Click Add Record, then select “Add Machine (A).”
- 17 Select the newly created newMachine.
- 18 Change the Machine Name field to the private hostname of the controller (for example, cluster).
- 19 Double-click the first IP address in the IP Address list and then change the first IP address to the public IP address for the controller (for example, 10.0.2.199).
- 20 Click Save.
- 21 Repeat steps 16 through 20 for each compute node using the private IP address reserved for them.

For example, the name of the first compute node is node1 assigned to 172.16.1.1, node2 assigned to 172.16.1.2, and so on.

- 22 Click the Start DNS button (below the Servers list).
The DNS service status indicator turns green when the service starts.
- 23 From the Apple Menu open System Preferences (/Applications/System Preferences).
- 24 Click Network.
- 25 Select the Ethernet 1 interface.
- 26 In the DNS Server field enter the public IP address of the controller (for example, 10.0.2.199).
- 27 In the Search Domains field enter the private DNS domain (for example, cluster).
- 28 Click Apply.
- 29 Quit System Preferences.

Verifying DNS Settings

Open Directory requires correct configuration of the DNS service. Before configuring the Open Directory Master, verify your DNS settings carefully. Any incomplete or incorrect Open Directory configuration prevents the cluster from functioning.

To verify DNS settings:

- 1 From the Dock on the cluster controller open the Terminal application.
- 2 Verify the fully qualified DNS name of the cluster controller using the `hostname` command.
For example, entering `hostname` returns `controller.cluster`.

```
$ hostname  
controller.cluster
```
- 3 Verify that the hostname of the cluster controller matches its assigned IP address in DNS using the `host` command.
For example, entering `host controller` returns `10.0.2.199`.

```
$ host controller  
controller.cluster has address 10.0.2.199
```
- 4 Verify that the fully-qualified DNS name of the cluster controller matches its public IP address using the `host` command.
For example, entering `host controller.cluster` returns `10.0.2.199`.

```
$ host controller.cluster  
controller.cluster has address 10.0.2.199
```
- 5 Verify that the reverse DNS record of the controller matches its fully-qualified DNS name using the `host` command.
For example, entering `host 10.0.2.199` returns `controller.cluster`.

```
$ host 10.0.2.199  
199.2.0.10.in-addr.arpa domain name pointer controller.cluster
```

If any DNS lookups do not match, repeat the process to create the DNS zone and entry for the controller. Do not continue the cluster setup process until DNS resolves correctly.

- 6 Quit Terminal.

Configuring Open Directory Service

The Open Directory service is responsible for authenticating users, publishing server setup configurations, and publishing network share automount records.

Configuring the Cluster Controller as an Open Directory Master

Use Server Admin to configure the Open Directory service on the cluster controller.

To configure Open Directory settings:

- 1 Open Server Admin if it is not already open.
- 2 In the controller's list of services, click Open Directory.
- 3 Click Settings, click General, then click Change.

This opens the Open Directory service configuration assistant.

- 4 Select Open Directory Master, then click Continue.
- 5 Create a Directory Administrator account, then click Continue.

Name, Short Name, User ID, Password: The Directory Administrator account administers the Open Directory domain that all nodes share. You can use the default Name, Short Name, and User ID. Choose a unique password.

- 6 Enter the Master Domain information, then click Continue.

Kerberos Realm: This field is preset to be the same as the server's private fully qualified DNS name converted to capital letters. Use the preset Kerberos Realm (for example, CONTROLLER.CLUSTER).

Search Base: This field is preset to a search base suffix for the new LDAP directory, derived from the private DNS name of the cluster controller. Use the preset LDAP search base (for example, dc=controller,dc=cluster).

WARNING: If these fields are not prepopulated, it might indicate your DNS settings were not configured properly. If so, click the Cancel button and redo the steps listed in "Configuring DNS Service" on page 84.

- 7 Confirm settings, then click Continue.
- 8 When the service configuration assistant completes, click Close.
- 9 Verify the Role is set to Open Directory Master.

Note: You can click Logs and monitor the log file /Library/Logs/slapconfig.log for errors while the Open Directory domain is being created. You can also use the Console (located in /Applications/Utilities/) or Terminal with the command “tail -f/Library/Logs/slapconfig.log.” In the log, warnings such as the following can be ignored:

```
WARNING: no policy specified for [...] defaulting to no policy
```

After the Open Directory domain is created, the Open Directory service starts and the status icon turns green.

Configuring DHCP Service

Using Server Admin, configure DHCP service on the cluster controller to provide LDAP and DNS information to the compute nodes.

To configure DHCP service:

- 1 Open Server Admin if it is not already open.
- 2 In the controller’s list of services, click DHCP.
- 3 Click Subnets.
- 4 Remove all subnets.
- 5 Create a new subnet for Ethernet port 2.
- 6 Click General and do the following:
 - a In the Subnet Name field, enter a subnet name (for example, Cluster Private Network).
 - b In the Starting IP Address field, enter the first IP address in the private network range available for compute nodes (for example, 172.16.1.1).
 - c In the Ending IP Address field, enter the last IP address in the private network range available for compute nodes (for example, 172.16.1.99).

Note: Leave some addresses unused at the end of the range for other devices and VPN connections.
 - d In the Subnet Mask field, enter the subnet mask for your private network (for example, 255.255.255.0).
 - e From the Network Interface pop-up menu, select en1 if it is not already selected. This menu shows the UNIX name for the port. The UNIX name for Ethernet 2 should be en1.
 - f In the Router field, enter the private IP address of the cluster controller (for example, 172.16.1.254).
 - g Set the lease time for the IP addresses served by the DHCP service to at least 1 month.
- 7 Click Save.

- 8 Click DNS below the Subnets list.
- 9 In the DNS Servers field, enter the public address of the cluster controller (for example, 10.0.2.199).
- 10 In the Default Search Domain field, enter the DNS domain for your private network (for example, cluster).
- 11 Click Save.
- 12 Click LDAP.
- 13 In the Server Name field, enter the fully qualified DNS name of the cluster controller (for example, controller.cluster).
- 14 In the Search Base field, enter the LDAP search base for your shared Open Directory domain (for example, dc=controller, dc=cluster).

This entry should match the LDAP search base entry you made when you created the Open Directory domain.

Note: Verify the Server Name and Search Base fields. Errors in the LDAP configuration of your DHCP service prevent proper autoconfiguration of cluster nodes, automounting of network directories, and use of network user accounts.

To avoid typographical errors, copy and paste the search base settings from the Open Directory service search base settings.
- 15 Select the Enable checkbox to the left of the subnet you just created.
- 16 Click Save.
- 17 Click the Start DHCP button (below the Servers list).

Configuring Firewall Settings on the Cluster Controller

The firewall on the controller is configured to enable access to all protocols from the public and private networks, but more limited access (for SSH and VPN) from external networks, including the Internet. You can adjust these rules to narrow or expand access to your controller.

To configure firewall settings on the cluster controller:

- 1 In the controller's list of services, click Firewall.
- 2 Click Settings, then click Address Groups.
- 3 From the IP Address Groups list, remove all entries except for "any."
- 4 Click the Add (+) button.
- 5 In the Group name field, enter the name of your public network (for example, example.com).
- 6 In the "Addresses in group" field, change the first entry to match your public IP network in CIDR notation.

For a subnet mask of 255.255.255.0, use "/24" after the network address (for example, 10.0.2.0/24).

- 7 Verify that the address range for the list accurately describes the address range used by your public network.
- 8 Click OK.
- 9 Click the Add (+) button to add another IP address group.
- 10 In the "Group name" field, name the group with your private DNS domain name (for example, cluster).
- 11 In the "Addresses in group" field, change the first entry to match your private IP network in CIDR notation.

For a subnet mask of 255.255.255.0, use "/24" after the network address (for example, 172.16.1.0/24).

- 12 Click OK.
- 13 Click Save.
- 14 Click Services.
- 15 From the "Edit Services for" pop-up menu, choose "any."
- 16 Select "Allow only traffic from 'any' to these ports."
- 17 Select the following ports (in addition to what's already selected):
 - ESP - Encapsulating Security Payload protocol
 - IKE NAT Traversal
 - VPN ISAKMP/IKE (500)
 - VPN PPTP—Point-to-Point Tunneling Protocol (1723)

Note: Enabling SSH and VPN ports on the controller allows remote access to the controller from your public network. Your public network can also be protected by a firewall service or device. If you plan to access your cluster from outside your public network (for example, using the Internet), talk to your system administrator about enabling the same ports on that firewall as well.

- 18 Click Save.
- 19 From the "Edit Services for" pop-up menu, choose the public network that was created in step 5 (for example, example.com).
- 20 Select "Allow all traffic from <public network>."
- 21 Click Save.
- 22 From the "Edit Services for" pop-up menu, choose the private network that was created in step 10 (for example, cluster).
- 23 Select "Allow all traffic from <private network>."
- 24 Click Save.

- 25 Click the Start Firewall button (below the Servers list).

Configuring NAT Settings on the Cluster Controller

Network Address Translation (NAT) allows compute nodes to share the controller's connection to the public network.

To configure NAT:

- 1 In the controller's list of services, click NAT.
- 2 Click Settings, then verify that IP Forwarding and Network Address Translation (NAT) is selected.
- 3 Verify that the "External network interface" pop-up menu is set to your public Ethernet interface (for example, Ethernet 1).
- 4 Verify that the Enable NAT Port Mapping Protocol checkbox is selected.
- 5 Click the Start NAT button (below the Servers list).

Configuring NFS

Using Server Admin, configure the NFS service on the cluster controller. NFS is used for file sharing and network home directory mounts.

To configure NFS service:

- 1 In the controller's list of services, click NFS.
- 2 Click Settings.
- 3 In the "Use__server threads" field, enter a number to specify the maximum number of NFS threads, or daemons, you want to run at one time.

An nfsd daemon is a server process that runs continuously behind the scenes and processes read and write requests from clients. The more threads that are available, the more concurrent clients can be served.

- 4 Click Save.
- 5 Click the Start NFS button (below the Servers list).

Configuring VPN Service

Configure the VPN service to enable secure connections from computers on remote networks.

To configure VPN service:

- 1 In the controller's list of services, click VPN.
- 2 Click Settings, then click PPTP.
- 3 Select the Enable PPTP checkbox.

- 4 In the Starting IP address field, enter the first private IP address you want to assign to remote VPN clients (for example, 172.16.1.200).
- 5 In the Ending IP address field, enter the last private IP address you want to assign to remote VPN clients (for example, 172.16.1.229).
- 6 Click Save.
- 7 Click the Start VPN button (below the Servers list).

Configuring Xgrid Service

Using Server Admin on the cluster controller, configure it as an Xgrid controller and then start Xgrid service.

Note: Because the cluster controller is also responsible for authentication, NFS sharing, network services, and possibly other critical services, it is not advisable for a cluster controller to run the Xgrid agent.

To configure the Xgrid service:

- 1 In the controller's list of services, click Xgrid.
- 2 Click Overview.
- 3 Click Configure Xgrid Service.
The service configuration assistant will launch.
- 4 Click Continue.
- 5 Select "Host a grid," then click Continue.
- 6 Enter the directory administrator's user name and password.
This is the directory administrator account you created when you enabled the Open Directory service.
- 7 Click Continue.
- 8 Verify that the Xgrid settings include the correct Kerberos realm (for example, CONTROLLER CLUSTER).
- 9 Click Continue.
- 10 Once the Xgrid service is configured, click Close.
- 11 Click Settings.
- 12 Click Agent, then deselect Enable Agent Service.
- 13 Click Save.
- 14 When prompted to restart Xgrid, click Restart.

Preparing the Data Drive as a Mirrored RAID set

When preparing your data drive you should protect your data by using a mirrored RAID set, also referred to as RAID 1. You can use the Disk Utility application to create the mirrored RAID set. To create a mirrored RAID set you must have two or more disks.

Note: Your network share points should be located on a different drive than your operating system, ideally on a mirrored RAID set.

To prepare the data drive as a mirrored RAID set:

- 1 Open the Disk Utility application (in /Applications/Utilitie).
- 2 From the drive list on the left, click one of the two drives to be used in the RAID.
- 3 Click RAID.
- 4 Enter a name for the RAID set (for example, Data).
- 5 Drag the disks you want to mirror from the left side of the pane to the disk list at the center of the pane.
- 6 For each disk you dragged to the disk list, verify the disk type is set to “Raid Slice.”
To use the disk as a mirror at all times, select RAID Slice.
To use the disk as a mirror only when another disk fails, select Spare.
- 7 To automatically rebuild mirror data, click Options, select “Automatically rebuild RAID mirror sets,” and then click OK.
- 8 Select the RAID set from the disk list and then from the Volume Format pop-up menu choose either “Mac OS Extended (Journaled)” or “Mac OS Extended (Case-sensitive, Journaled).”
If you plan to work with applications or source code that was designed for other UNIX operating systems, choose the case-sensitive option.
- 9 From the RAID Type pop-up menu, choose Mirrored RAID Set.
- 10 Click Create.
- 11 Select the mirrored RAID that will host your data volume.
- 12 Use the cluster administrator username and password to authenticate.
- 13 Verify that the RAID set has the correct format.
- 14 Quit the Disk Utility application.

Creating a Home Directory Automount Share Point

Use Server Admin to configure an automount share point on the cluster controller.

To create an automount home directory share point:

- 1 Open Server Admin and select the controller in the Servers list.
- 2 Click File Sharing, then click Volumes.
- 3 Select the volume you want to contain the home directory share point (for example, Data).
- 4 Click Browse.
- 5 Click New Folder, name the folder "home," then click Create.
- 6 Click Save.
- 7 Select the home folder you created.
- 8 Click Share, then click Share Point.
- 9 Select Enable Automount.

The Automount configuration screen appears.

- 10 Verify that the directory is set to /LDAPv3/127.0.0.1.
- 11 From the protocol pop-up menu choose NFS.
- 12 Verify that "Use for" is set to User home folders.
- 13 Click OK.
- 14 When prompted, enter the directory administrator's user name and password.
- 15 Deselect "Enable Spotlight searching."
- 16 From Share Point, click Protocol Options.

The Protocol Options screen appears.

- 17 Click NFS.
- 18 Select the "Export this item and its contents to" checkbox, then choose Subnet from the pop-up menu.
- 19 Set the Subnet address field to your private network address (for example, 172.16.1.0).
- 20 Set the Subnet mask field to your private network subnet mask (for example, 255.255.255.0).
- 21 Verify that the mapping pop-up menu is set to "Root to Nobody."
- 22 Click OK.
- 23 Click Save.
- 24 Restart the controller (Apple Menu > Restart).

Creating User Accounts

Use Workgroup Manager to create user accounts.

To create user accounts:

- 1 If you did not restart the cluster controller at the end of the previous section (“Creating a Home Directory Automount Share Point” on page 93), restart it now.
- 2 Log in using your administrator account.
- 3 Open Workgroup Manager (located at /Applications/Server/).
You can also open Workgroup Manager from the Dock.
- 4 Connect to the cluster controller using its hostname and your administrator user name and password.
- 5 On the right side of the Workgroup Manager window, click the lock button.
- 6 Authenticate with the directory administrator username and password.
- 7 Click Accounts.
- 8 Select the users icon tab above the accounts listing.
- 9 Click New User.
- 10 In the Name field, enter the full name for a user (for example, “Tom C”).
- 11 In the Short Names list box, enter a short username for the user (for example, “tac”).
- 12 In the Password field, enter a password for the user.
- 13 In the Verify field, reenter the password for the user.
- 14 Click Save.
- 15 Click Advanced.
- 16 From the Login Shell pop-up menu, choose the preferred shell for the user.
- 17 Click Home.
- 18 From the list, select the NFS automount share point (home).
- 19 Click Create Home Now.
- 20 Click Save.
- 21 Repeat this process for each cluster user.
- 22 Quit Workgroup Manager.

Simplify the compute node setup process by creating Auto Server Setup records.

An Auto Server Setup record is an XML property list with values that can be used to automatically complete the Server Assistant for newly installed Mac OS X servers. Auto Server Setup records can be accessed using external storage (for example a CD, USB drive, or iPod) or over a network using Open Directory.

For more information about creating and using Auto Server Setup records, see *Server Administration*.

You can accomplish additional automation of compute node configuration by using scripts executed with SSH or Apple Remote Desktop software.

Creating an Auto Server Setup Record for Compute Nodes

To automate the process of setting up compute nodes, use Server Assistant to save the compute node configuration to a file or Open Directory record.

To create an Auto Server Setup record:

- 1 On the cluster controller, open Server Assistant (located in `/Applications/Server/`).
- 2 In the Welcome screen:
 - a Select “Save advanced setup information in a file or directory record.”
 - b Click Continue.
- 3 In the Language screen:
 - a Select the language you want to use to administer the server.
 - b Click Continue.
- 4 In the Keyboard screen:
 - a Select the keyboard layout for the server.
 - b Click Continue.

- 5 In the Serial Number screen:
 - a Enter a site-licensed Mac OS X Server serial number.
Note: If you don't have a site-licensed number you must manually enter unique serial numbers for each compute node after it has been configured.
 - b Click Continue.
- 6 In the Administrator Account screen:
 - a Create the account you'll use to administer compute nodes.
 - b Click Continue.
- 7 In the Network Interfaces screen:
 - a Click Add.
 - b In the Port Name field, enter "Ethernet 1."
 - c In the Device Name field, enter "en0" and leave the Ethernet Address field blank.
 - d Click OK.
 - e Enable TCP/IP for Ethernet 1.
 - f Click Continue.
- 8 In the TCP/IP Connection screen for the Built-in Ethernet 1 port:
 - a From the Configure pop-up menu, choose Using DHCP.
 - b Leave the other fields blank.
 - c Click Continue.
- 9 In the Network Names screen:
 - a Leave the Primary DNS Name field blank.
 - b Leave the Computer Name field blank.
 - c Verify that the "Enable Remote Management" checkbox is selected.
 - d Click Continue.
A warning appears indicating you left some fields blank.
 - e Click Continue.
- 10 In the Time Zone screen:
 - a From the Closest City pop-up menu, choose your time zone.
 - b Click Continue.

- 11 In the Directory Usage screen:
 - a From the “Set directory usage to” pop-up menu, choose “Connected to a Directory System”.
 - b From the Connect pop-up menu, choose “Open Directory Server.”
 - c In the IP Address or DNS Name field, enter the private DNS name of the cluster controller (for example, controller.cluster).
 - d Click Continue.
- 12 In the Confirm Settings screen:
 - a Read the configuration summary to confirm that you have made the correct settings.
 - b Click Save As.
- 13 In Save settings, use the following to choose whether to save your setting to a configuration file or Open Directory record.

If you use a configuration file, it should be named generic.plist and saved to a CD, DVD, USB drive, iPod, or other removable drive. It should be located in a folder called Auto Server Setup at the top level of the removable file system. The file is used if the removable drive is present when an unconfigured compute node starts for the first time.

If you save your settings to an Open Directory record, an unconfigured compute node discovers the record via DHCP and configures itself accordingly. Save the record to the LDAPv3/127.0.0.1 domain and name it generic. When asked, specify an Open Directory server using the controller’s DNS name (for example, controller.cluster) or IP address (for example, 10.0.2.199).

Saving settings to an Open Directory record without encryption will require the use of password (.pass) files. Saving them without encryption will expose the administrator password to anyone with access to the Open Directory domain. For more information about the creation and use of Auto Server Setup record and encryption, see *Server Administration*.

 - a Select Directory Record.
 - b If creating a Directory Record, choose /LDAPv3/127.0.0.1 from the Directory Domain pop-up menu.
 - c Decide if you want to encrypt the record.
 - d In the Record Name field, enter “generic.”
 - e Click OK and then authenticate using the directory admin login and password you created when you configured Open Directory.
- 14 Click OK.
- 15 Quit Server Assistant.

Verifying LDAP Record Creation

To verify the creation of the LDAP directory record that will be used by compute nodes to autoconfigure, use the `slapcat` command on the cluster controller.

To verify the LDAP record creation:

- 1 Open a Terminal window on the cluster controller and enter the following command:

```
$ sudo slapcat | grep generic
```

- 2 When prompted enter the administrator password .

This command displays the generic records in the LDAP database on the cluster controller. In this case, there should only be one record—the one you created in the previous section.

```
dn: cn=generic,cn=autoserverssetup,dc=controller,dc=cluster
cn: generic
```

Setting Up Compute Nodes

Setting up compute nodes involves obtaining IP addresses for each compute node connected to your private network. This section provides useful tips for setting up compute nodes depending on your cluster configuration.

To set up compute nodes:

- 1 Make sure compute nodes are connected to the private network through Ethernet port 1.
- 2 Start the first compute node.

The DHCP service hosted on the cluster controller provides IP addresses to nodes when they start, beginning with the first address in the range and incrementing the address for each request. The DHCP lease time specified in the Server Admin settings for the DHCP service determines how long this address is reserved for a computer.

It is advisable for each node in a cluster to use sequential IP addresses that correspond to their physical position in a rack and the names they have been assigned. Node1 would have an address that ends in 1 (for example, 172.16.1.1) and node199 would have an address that ends with 199 (for example, 172.16.1.199).

If you set up your cluster in this manner, start the first node and wait until you verify its IP address before starting the next one. You can check DHCP IP address assignments in the DHCP Clients pane of Server Admin. Because Server Admin does not maintain a persistent connection to the servers it administers, you might need to click the Refresh button in the toolbar to update the client listing immediately.

If an Auto Server Setup record is available to the compute node through a removable drive or Open Directory record, it will configure itself and reboot. After you verify that the first node has completed this process, start the remaining compute nodes sequentially, allowing time for them to obtain sequential IP addresses from the DHCP server and for autoconfiguration. Do not disconnect or remove disks until you are sure the server has applied the settings.

- 3 Select the DHCP service and view client connections.

Static Maps in the DHCP Static Maps pane of Server Admin enable you to guarantee that an IP address is always reserved for a specific node, regardless of how much time has elapsed since it was assigned its address.

In addition to providing the IP address assignment, the DHCP service on the cluster controller provides the IP address and search base for the Open Directory domain on the cluster controller.

Configuring Cluster Nodes

When configuring cluster nodes, use Server Admin to name cluster nodes, join them to the Kerberos realm, and join them to a grid.

To configure cluster nodes:

- 1 Open Server Admin.
- 2 Click the Add Server (+) button below the Servers list.
- 3 Connect to the cluster node using its IP address.

If you used an Auto Server Setup record to configure the nodes, use the administrator user name and password you created with that record.

- 4 In the Servers list, click the cluster node.
- 5 Click Settings.

Note: If the Mac OS X Server serial number is not valid, Server Admin doesn't permit you to administer services. If you did not supply a volume license serial number when creating the Auto Server Setup file, you must enter a valid serial number for each node before you can continue. Click General to verify the serial number.

- 6 Click Network.
- 7 In the Computer Name and Local Hostname fields, enter the computer name and hostname of the cluster node (for example, node1).
- 8 Click Save.
- 9 Click Services.
- 10 Select the Open Directory checkbox.
- 11 Select the Xgrid checkbox.

- 12 Click Save.
- 13 Repeat steps 2 through 12 for each compute node.

You can also use Apple Remote Desktop to set the names of all cluster nodes at once. For more information, see “Naming Multiple Cluster Nodes” on page 111.
- 14 Select the node’s Open Directory service.
- 15 Click Settings, then click General.
- 16 Verify the role is set to “Connected to a Directory System.”
- 17 Click Join Kerberos.

A Join Kerberos Realm screen appears. Set the realm to your Kerberos realm (for example, CONTROLLER.CLUSTER).
- 18 Enter the Open Directory administrator user name and password.
- 19 Click Refresh below the Servers list.

If the node has joined the Kerberos realm, the Join Kerberos button and associated text will disappear.
- 20 In the Servers list select the node’s Xgrid service.
- 21 Click Overview.
- 22 Click Configure Xgrid Service.

The Xgrid Service Configuration Assistant appears.
- 23 Click Continue, then select “Join a grid.”
- 24 Click Continue.
- 25 In the “Use controller with hostname” field, enter the controller’s private DNS name (for example, controller.cluster).
- 26 Click Continue.
- 27 Confirm the settings.

The Directory Server entry should be an LDAPv3 path based on the controller’s DNS name (for example, /LDAPv3/controller.cluster). The Kerberos realm should be the same as the controller’s DNS name in all capital letters (for example, CONTROLLER.CLUSTER).
- 28 Click Continue.
- 29 Click Close.

You can automate steps. For more information, see Appendix B, “Automating Compute Node Configuration.”

Creating and Verifying a VPN Connection

Remote clients can connect to the private network of the cluster securely using SSH and VPN. VPN access allows graphical applications (like the GridMandelbrot sample Xgrid application) to run on remote systems, but use the cluster for computation. VPN access also allows administrative tools, such as Apple Remote Desktop, to manage compute nodes from a remote system.

The following instructions are for VPN configuration for Mac OS X v10.5 clients. For other operating systems, or older versions of Mac OS X, consult the appropriate documentation using the values provided in the following.

To create and verify a VPN connection:

- 1 Open System Preferences, then click Network.
- 2 Click the Add (+) button at the bottom of the network connection services list and then choose VPN from the Interface pop-up menu.
- 3 From the VPN Type pop-up menu, choose PPTP.
- 4 In the Service Name field, enter a descriptive name (for example, Cluster VPN)) and click Create.
- 5 In the Server Address field, enter the public IP address for the controller (for example, 10.0.2.199).
- 6 In the Account Name field, enter the short username for a user you created on the controller using Workgroup Manager.
For more information, see “Creating User Accounts” on page 94.
- 7 Click Apply and then click Connect.
- 8 Verify that the network connection services list has an active VPN (PPTP) connection to the cluster controller and that you’re getting a private network address.

Joining a Remote Client to the Kerberos Realm

Because the firewall has been configured to block most types of incoming network access, a VPN connection is necessary to use Kerberos from remote clients. For your client computer to use Kerberos, you must join it to the Kerberos realm of the controller.

To join a remote client to the Kerberos realm:

- 1 Open the Kerberos application located in the /System/Library/CoreServices/ folder.
- 2 Select Edit > Edit Realms.
- 3 Click the Add (+) button below the Realm list.
- 4 In the Realm Name field, enter the Kerberos Realm of the controller (for example, CONTROLLER.CLUSTER).

- 5 Click Servers, then click the Add (+) button (below the Servers list).
- 6 Verify that the new entry in the Type column is listed as "KDC."
- 7 Enter the private DNS name for your controller in the Server column (for example, controller.cluster).
- 8 Click Domain, then click the Add (+) button (below the Domain list).
- 9 Enter the private DNS zone preceded by a period (for example, .cluster).
- 10 Click the Add (+) button (below the Domain list).
- 11 Enter the private DNS zone (for example, cluster).
- 12 Click OK.
- 13 Authenticate using administrator credentials for you client computer.

Verifying Remote Client Access to the Kerberos Realm

After the remote client is configured to join the Kerberos realm, verify that you have received a Ticket Granting Ticket (TGT) from the controller.

To verify remote client access to the Kerberos realm:

- 1 Open the Kerberos application located in the /System/Library/CoreServices/ folder.
- 2 Click New.
- 3 Verify that the Realm is set to the Kerberos Realm of the controller (for example, CONTROLLER.CLUSTER).
- 4 Enter the user name and password for an account created in the Open Directory domain of the controller.
- 5 Click OK.
- 6 Verify the entry in the Ticket Cache list.
- 7 Verify the entry of the TGT for your user in the Ticket list (for example, krbtgt/CONTROLLER.CLUSTER@CONTROLLER.CLUSTER).

Note: When an application that supports Kerberos is used and the Kerberos TGT does not exist or has expired, the Kerberos authentication dialog appears. You do not need to use the Kerberos application each time you want to obtain a ticket.

Use this chapter to make sure you've successfully configured your cluster before performing HPC.

Use Xgrid Admin to verify that you can see the Xgrid agents in your cluster. Then use sample Xgrid tasks to test your cluster.

Checking Your Cluster Using Xgrid Admin

Use Xgrid Admin to verify that Xgrid agents are running on the compute nodes.

To use Xgrid Admin to check your cluster:

- 1 From the management computer, a VPN client, or the controller, open Xgrid Admin (located in /Applications/Server/).
- 2 Click Add Controller.
- 3 From the pop-up menu, choose the controller and click Connect.
- 4 In the authentication sheet:
 - a Select "Use Single Sign On Credentials."
 - b Click OK.
 - c If prompted, enter a cluster account username, the Kerberos realm (for example, CONTROLLER.CLUSTER), and password.
 - d Click OK.
- 5 In the Controllers and Grids list, select the cluster.
- 6 Click Overview.

Overview shows the number of agents, which should equal the number of compute nodes you configured.

This also shows the number of available, unavailable, and working processors, and the number of jobs running and jobs pending.
- 7 View the status of the Xgrid agents by clicking Agents.

- 8 Verify that you can see a list of all nodes in your cluster.
If you don't see all agents you were expecting, see "If Your Agents Can't Connect to the Xgrid Controller" on page 51.
- 9 Monitor the progress of Xgrid jobs as they are being processed by clicking Jobs.
- 10 Quit Xgrid Admin.

Testing Your Xgrid Cluster

To test your cluster, use GridSample, a sample Cocoa application that comes with Developer Tools for Mac OS X v10.5, to submit Xgrid tasks to the controller. This application provides you with an easy-to-use GUI for Xgrid. On any system that has the Mac OS X developer tools installed, the example code for the application is at:

```
/Developer/Examples/Xgrid/GridSample/GridSample.xcodeproj
```

Using this application, you can generate the monthly calendars of the year 2007 across the cluster. Although this application is trivial, it enables you to test the cluster and it illustrates the simplicity of Xgrid job submission.

Note: You can also submit Xgrid tasks using the `xgrid` command-line tool. For more information, see the tool's man page and *Command-Line Administration*.

To test your cluster using GridSample:

- 1 Open GridSample.xcodeproj by using Xcode (located in /Developer/Applications/).
- 2 Set the active executable to Xgrid Feeder Sample by choosing Project > Set Active Executable > Xgrid Feeder Sample.
- 3 Build and run the project by clicking "Build and Go."
The application starts running and prompts you for an Xgrid controller to connect to.
- 4 Enter the address of the controller and click Connect.
- 5 Click "Use password," enter the password for the controller, and click OK.
- 6 Click New Job.
- 7 In the Job Name field, enter "2007 Calendars."
- 8 Make sure the Command field is set to /usr/bin/cal.
- 9 From the Argument 1 pop-up menu, choose Range.
- 10 For argument 1, enter 1 in the From field, 12 in the "to" field, and 1 in the "by" field.
This range tells the application to generate the 2007 monthly calendars from January through December.
- 11 To add another argument below Argument 1, click the Add (+) button.
- 12 From the Argument 2 pop-up menu, choose Literal.

- 13 For argument 2, enter "2007."

Note: Instead of specifying one year, you could specify a range of years, and Xgrid would create a separate set of tasks for each year.

- 14 Click Submit.

The Xgrid controller on the controller prepares the tasks and sends them to Xgrid agents running on the cluster nodes. When the job is done, the status of the job changes to Finished in the Xgrid Feeder Sample window.

- 15 To see the results of each task, click Show Results.

Note: To test image-rendering on your cluster, use Xcode to build and run the example application GridMandelbrot.xcodeproj (located in /Developer/Examples/Xgrid/GridMandelbrot/). Just as you did earlier, build and run the project, connect to the Xgrid controller, and submit the job. The application renders Mandelbrot images across your cluster.

Verifying Your Xgrid Configuration

Verify that Xgrid is configured and works.

To verify your Xgrid service:

- 1 Install and configure Xcode developer tools.
Xcode is included with the Mac OS X Server Installation disc. The latest version of Xcode can also be downloaded from the Apple Developer Connection (ADC) at www.apple.com/developer.
- 2 Compile and launch the Xgrid Mandelbrot example application (located in /Developer/Examples/Xgrid/GridMandelbrot).
- 3 From the "Enter or choose a controller to connect to" pop-up menu, choose your controller and click Connect.
- 4 Select "Use Single Sign On credentials" and click OK.
- 5 Enter a cluster user name and password to authenticate with Kerberos, then click OK.

You can monitor your cluster's performance with the Xgrid Admin application in /Application/Server/.

Verifying Your SSH Connection

Verify that SSH is running on the controller by using Terminal.

To verify your SSH connection:

- 1 From a remote system, open Terminal (located in /Applications/Utilities/).
- 2 Open an SSH connection to your controller by logging in with a user account name and password created in Workgroup Manager and by using the public IP address or public DNS name for your controller (for example, `ssh tomclark@10.0.2.199`).

Enter the following command to obtain a Kerberos Ticket Granting Ticket (TGT) and when prompted for a password use the same password used for your SSH connection.

By using a TGT you are not required to enter passwords for access to cluster resources.

```
$ kinit
```

```
Please enter the password for tomclark@CONTROLLER.CLUSTER:
```

After the connection to the controller is made, you can connect directly to the compute nodes using their private DNS name (for example, `ssh tomclark@node1.cluster` OR `ssh tomclark@node1`).

Cluster Setup Checklist

A

Use the checklist in this appendix to guide you through the cluster setup procedure.

Print this checklist and use it to make sure you have performed all setup steps. The steps in this checklist are in order only within each section.

For information about this step	Go to
Physical Setup	
<input type="checkbox"/> Power source meets minimum requirements	"Power Requirements" on page 68
<input type="checkbox"/> Cooling system meets minimum requirements	"Cooling Requirements" on page 69
<input type="checkbox"/> Facility housing the cluster meets minimum weight requirements	"Weight Requirements" on page 70
<input type="checkbox"/> Space around the cluster meets minimum requirements	"Space Requirements" on page 70
<input type="checkbox"/> Network switches support Gigabit Ethernet and have enough ports	"Network Access Requirements" on page 71
<input type="checkbox"/> Mount cluster nodes on the rack	"Network Access Requirements" on page 71
<input type="checkbox"/> Connect cluster nodes to a power source	"Preparing the Cluster Nodes for Software Configuration" on page 75
<input type="checkbox"/> Connect cluster nodes to the private network	"Preparing the Cluster Nodes for Software Configuration" on page 75
Software Setup	
<input type="checkbox"/> Obtain a static IP address and related network and DNS information	"Network Access Requirements" on page 71
<input type="checkbox"/> Obtain a site-licensed serial number	"Volume-Licensed Serial Number" on page 72
<input type="checkbox"/> Obtain a copy of Apple Remote Desktop	"Apple Remote Desktop" on page 72
<input type="checkbox"/> Record the serial numbers of cluster nodes	"Preparing the Cluster Nodes for Software Configuration" on page 75

For information about this step	Go to
Management Computer Setup (Optional)	
<input type="checkbox"/> Disable AirPort and other public network connections	"(Optional) Setting Up the Management Computer" on page 78
<input type="checkbox"/> Install the latest version of Mac OS X Server tools	"(Optional) Setting Up the Management Computer" on page 78
<input type="checkbox"/> Install Apple Remote Desktop	"(Optional) Setting Up the Management Computer" on page 78
Controller Setup	
<input type="checkbox"/> Connect the controller to the public and private network	"Setting Up Server Software on the Cluster Controller" on page 81
<input type="checkbox"/> Run Server Assistant and configure public network settings	"Setting Up Server Software on the Cluster Controller" on page 81
<input type="checkbox"/> Configure DNS service	"Configuring DNS Service" on page 84
<input type="checkbox"/> Configure Open Directory service	"Configuring the Cluster Controller as an Open Directory Master" on page 86
<input type="checkbox"/> Configure DHCP service	"Configuring DHCP Service" on page 87
<input type="checkbox"/> Configure Firewall service	"Configuring Firewall Settings on the Cluster Controller" on page 88
<input type="checkbox"/> Configure NAT service	"Configuring NAT Settings on the Cluster Controller" on page 90
<input type="checkbox"/> Configure NFS service	"Configuring NFS" on page 90
<input type="checkbox"/> Configure VPN service	"Configuring VPN Service" on page 90
<input type="checkbox"/> Configure Xgrid service	"Configuring Xgrid Service" on page 91
<input type="checkbox"/> Prepare data drive	"Preparing the Data Drive as a Mirrored RAID set" on page 92
<input type="checkbox"/> Create home directory	"Creating a Home Directory Automount Share Point" on page 93
<input type="checkbox"/> Create user accounts	"Creating User Accounts" on page 94
Compute Node Setup	
<input type="checkbox"/> Create auto server setup records	"Creating an Auto Server Setup Record for Compute Nodes" on page 95
<input type="checkbox"/> Set up compute nodes	"Setting Up Compute Nodes" on page 98
<input type="checkbox"/> Configure cluster nodes	"Configuring Cluster Nodes" on page 99
<input type="checkbox"/> Create and verify VPN connection	"Creating and Verifying a VPN Connection" on page 101
Cluster Testing	
<input type="checkbox"/> Check the cluster using Xgrid Admin	"Checking Your Cluster Using Xgrid Admin" on page 103
<input type="checkbox"/> Test Xgrid cluster	"Testing Your Xgrid Cluster" on page 104

For information about this step	Go to
<input type="checkbox"/> Verify Xgrid configuration	"Verifying Your Xgrid Configuration" on page 105
<input type="checkbox"/> Verify your SSH connection	"Verifying Your SSH Connection" on page 106

Automating Compute Node Configuration

B

Use this appendix to learn about alternative ways of completing tasks documented earlier in this guide.

For large clusters, some tasks in this guide can be completed quickly and efficiently using Apple Remote Desktop.

Naming Multiple Cluster Nodes

Using the Send UNIX Command in Apple Remote Desktop, you can rename all cluster nodes at once.

The shell script used in the following steps causes each node to set its Computer name and Bonjour name to “node” followed by the last digit of its IP address. For example, a node with the IP address of “172.16.1.2” will be named “node2.”

To name multiple cluster nodes:

- 1 Open Apple Remote Desktop.
- 2 Select the nodes to be configured.
- 3 From the Manage pop-up menu, select “Send UNIX Command.”
- 4 In the first field, enter the following shell script, noting the use of double quotes (“) and backquotes (`).

```
theNodeNumber=`ipconfig getifaddr en0 | cut -d . -f 4`  
/System/Library/ServerSetup/serversetup -setComputerName  
    "node${theNodeNumber}"  
/System/Library/ServerSetup/serversetup -setBonjourName "node${theNodeNumber}"
```

- 5 Select button next to User.
- 6 In the User field, enter “root.”
- 7 Click Send.

For each node that sets its name, an entry is created in the results window followed by two lines containing a zero.

- 8 Close the Send UNIX Command results window.

All nodes should now show their hostname in the Remote Desktop list.

Joining Multiple Cluster Nodes to the Kerberos Realm

To send commands to join the nodes to the Kerberos realm, use Apple Remote Desktop's Send UNIX Command.

To join multiple cluster nodes to the Kerberos realm:

- 1 Open Apple Remote Desktop.
- 2 Select the nodes you want to join.
- 3 From the Manage pop-up menu, choose Send UNIX Command.
- 4 In the first field, enter the following command:

```
sso_util configure -r CONTROLLER.CLUSTER -a diradmin -p diradminpassword all
```

This command sets each cluster node to join the Kerberos realm "CONTROLLER.CLUSTER" using the directory administrator account "diradmin" and the password "diradminpassword."

- 5 Select the button next to User.
- 6 In the User field, enter "root".
- 7 Click Send.

For each node joining the Kerberos realm, there is an entry in the results window.

- 8 Close the Send UNIX Command results window.

Configuring Xgrid Agent Settings Using Apple Remote Desktop

To send commands to compute nodes to configure their Xgrid agent settings, use Apple Remote Desktop's Send UNIX Command.

To configure Xgrid agent settings:

- 1 Open Apple Remote Desktop.
- 2 From the pop-up menu, click Scanner and choose Network Range.
- 3 Enter the starting and ending addresses of the address range used by the compute nodes.
- 4 Select the compute nodes from the list and choose Manage > Send UNIX Command.

5 In the text field, enter the following commands:

```
serveradmin settings xgrid:XgridKerberosInfo:ReadyForAgentRoleBasedSetup =
    yes
serveradmin settings
    xgrid:XgridKerberosInfo:ReadyForControllerRoleBasedSetup = yes
serveradmin settings xgrid:AgentSettings:Enabled = yes
serveradmin settings xgrid:AgentSettings:ControllerPassword = ""
serveradmin settings xgrid:AgentSettings:prefs:ControllerName = "controller"
serveradmin settings xgrid:AgentSettings:prefs:SuspendWhenNotIdle = no
serveradmin settings xgrid:AgentSettings:prefs:OnlyWhenIdle = no
serveradmin settings xgrid:AgentSettings:prefs:ResolveNameAsNetService = yes
serveradmin settings xgrid:AgentSettings:prefs:ControllerAuthentication =
    "Kerberos"
serveradmin settings xgrid:AgentSettings:prefs:BindToFirstAvailable = no
serveradmin settings xgrid:ControllerSettings:ClientPassword = ""
serveradmin settings xgrid:ControllerSettings:Enabled = no
serveradmin settings xgrid:ControllerSettings:prefs:AgentAuthentication =
    "Kerberos"
serveradmin settings xgrid:ControllerSettings:prefs:ClientAuthentication =
    "Kerberos"
serveradmin settings xgrid:ControllerSettings:AgentPassword = ""
xgridctl agent start
```

Replace "controller" with the fully qualified private name of the controller (for example, controller.cluster).

- 6** Select User and enter "root" in the text field.
- 7** Select "Display all output."
- 8** Click Send.

These commands configure the Xgrid agent on compute nodes to bind to the controller and then start the Xgrid service.

The compute nodes can now receive Xgrid tasks.

Using SSH Without Passwords

Users on your cluster can generate authentication keys in their home folders that enable them to use SSH to connect to other cluster nodes without entering their password again.

To use SSH without passwords:

- 1 Make an SSH connection to the controller.

If connecting from a remote system, access the public IP address or DNS name of the controller (For example, `ssh mab@10.0.2.199`).

- 2 In your home directory on the controller, enter the following commands in sequence:

```
mkdir .ssh
chmod 700 .ssh
ssh-keygen -t dsa -f .ssh/id_dsa -C "Enter a comment here"
```

You are prompted twice to enter a passphrase. Leave this blank and press Return each time.

```
chmod 600 .ssh/id_dsa*
cat .ssh/id_dsa.pub >> .ssh/authorized_keys
```

You can test the authentication keys by attempting to make an SSH connection from the controller to a cluster node (for example, `ssh mab@node2.cluster`).

The first time you connect to any cluster node, SSH prompts you to establish the authenticity of that node by entering “yes” at the prompt. After the authenticity of the node is established, a record is stored in the `~/.ssh/known_hosts` file of your home folder and you are not prompted for that host again.

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

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 to share files and network services. AFP uses TCP/IP and other protocols to support communication 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 the purpose of recovery in case the original copy of data is lost or becomes inaccessible.

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 interacting with the computer (for example, to run programs or modify file system permissions) by entering text commands at a shell prompt. See also **shell**; **shell prompt**.

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.

data rate The amount of information transmitted per second.

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

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 drive A device that contains a disk and reads and writes data to the disk.

disk image A file that, when opened, creates an icon on a Mac OS X 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**.

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**.

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.

host name A unique name for a computer, historically referred to as the UNIX hostname.

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

Hypertext Transfer Protocol See **HTTP**.

image See **disk image**.

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

Internet 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 data packets and TCP keeps track of data packets.

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

KB Kilobyte. 1,024 (2^{10}) 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) To start a session with a computer (often by authenticating as a user with an account on the computer) 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 computer.

logical disk A storage device that appears to a user as a single disk for storing files, even though it might actually consist of more than one physical disk drive. An Xsan volume, for example, is a logical disk that behaves like a single disk even though it consists of multiple storage pools that are, in turn, made up of multiple LUNs, each of which contains multiple disk drives. See also **physical disk**.

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.

megabyte See **MB**.

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 can export shared volumes to computers based on IP address, and also supports single sign-on (SSO) authentication through Kerberos.

nfsd daemon An NFS server process that runs continuously behind the scenes and processes NFS protocol and mount protocol requests from clients. nfsd can have multiple threads. The more NFS server threads, the better concurrency.

NIC Network interface card. An adapter that connects a computer or other device to a network. NIC is usually used to refer to adapters in Ethernet networking; in Fibre Channel networking, the interface is usually called a host bus adapter (HBA).

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

open 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. 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.

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 See **SMB**.

SMB Server Message Block. 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. SMB services use SMB 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.

A

- access
 - administrator permissions 36
 - LDAP 86, 98
 - managing client 35
- accounts 94
- ACLs (access control lists) 35
- administrator 36, 42
- agents
 - adding 43
 - authentication 26
 - controllers 23, 30, 32, 91
 - deleting 44
 - distributed grids 21
 - functions of 22
 - grid workload 19
 - list of 43
 - management of 42, 43
 - mobility of 39
 - overview 23
 - requirements 18
 - setup 32, 33, 42, 43, 112
 - troubleshooting 51
- airflow for hardware 77
- Apple Remote Desktop (ARD)
 - agent settings 112
 - clusters 72
 - features 42
- Apple Workgroup Cluster 62
- applications
 - grid performance 19
 - Xgrid support 53
 - Xserve support 61
 - See also specific applications*
- ARD. *See* Apple Remote Desktop
- authentication
 - cluster 86, 112, 114
 - options 26, 27, 31
 - passwords 26, 27, 30, 34
 - setup 33, 34
 - troubleshooting 54
 - See also* Kerberos

B

- bioinformatics 62

C

- cat tool 37
- client computers, agent setup 33
- clients
 - access control 35
 - authentication 27
 - overview 23
 - remote, joining to Kerberos 101
 - verifying remote access 102
 - See also* client computers; users
- clusters
 - authentication 86, 112, 114
 - checklist for setup 107
 - connections 71
 - controllers 81, 91, 98
 - data drives 92
 - definition 20
 - DHCP 87, 98
 - DNS 84
 - domain for 99
 - high performance computing 59, 60
 - homogeneity of 20
 - management computer 78, 81
 - NFS 90, 93
 - Open Directory 86
 - requirements 67, 68, 73
 - setup overview 63, 64
 - testing 103, 104, 105, 106
 - user accounts 94
 - VPN 90
 - vs. grids 18
 - Xgrid capacity 24
 - Xgrid service 91
 - Xserve 60, 62
 - See also* nodes
- command-line tools
 - Server Admin 38
 - SSH login 42
 - viewing logs 37

- Xgrid 42, 48, 104
- computational grids. *See* grids, computational
- computers
 - agent setup 42
 - client 33
 - idle status 32
 - management 78, 81
- configuration
 - agents 32, 33, 42, 43, 112
 - authentication 33, 34, 86, 112
 - automatic grid 22
 - controller 30, 81, 91
 - hosting 28
 - joining 29
 - remote preferences 42
 - Service Configuration Assistant 28
 - See also* clusters; nodes
- controllers
 - and agents 23, 30, 32, 91
 - cluster 81, 91, 98
 - connections 40, 41
 - hosting considerations 28
 - management of 40
 - NAT settings 90
 - nodes 22
 - overview 24
 - requirements 18
 - security 88
 - setup 30, 81, 91
- cooling requirements 69
- cross-platform Xgrid agents 53
- cryptography 62

D

- data drive setup 92
- data mining 62
- desktop recovery 18
- DHCP (Dynamic Host Configuration Protocol)
 - service 87, 98
- directory services 84, 86, 99
- disk images 43
- disks, cluster preparation 92
- Disk Utility 92
- distributed computing architecture 21
 - See also* Xgrid
- DNS (Domain Name System) service 28, 84, 85
- documentation 11, 12, 13
- Domain Name System. *See* DNS
- domains, directory 84, 86, 99
- drives. *See* disks
- Dynamic Host Configuration Protocol. *See* DHCP

E

- embarrassingly parallel computations 62
- Ethernet

- Gigabit Ethernet 62, 71
- ports for 77, 81

F

- failure rates 21, 48
- file services 90, 93
- firewall service 28, 51, 88

G

- Gigabit Ethernet 62, 71
- grids, computational
 - automatic configuration 22
 - definition 18, 39
 - functions 22
 - management of 39, 45, 46
 - overview 17
 - performance 19
 - types 21
 - vs. clusters 18
 - See also* Xgrid
- GridSample 104

H

- hardware requirements 67, 68, 69, 70
- head node 21
- help, using 10
- highly dispersed grids 18
- high performance computing (HPC)
 - Apple's role in 59, 60, 62
 - overview 59
- homogeneity of clusters 20
- host names 73, 86
- HPC. *See* high performance computing

I

- images
 - disk 43
 - rendering of 62, 105
- indicators, status 40
- installation
 - NetInstall 43
- IP addresses
 - DHCP setup 87, 98
 - DNS service 84
 - hosting controller 28
 - static 73
 - VPN setup 91

J

- jobs
 - definition 24
 - deleting 45
 - failure of 48
 - list of 44
 - overview 18, 19, 22, 23, 24

- restarting 45
- results 49
- status checking 49
- stopping 44
- structuring 47
- styles 47
- submitting 48

K

- Kerberos
 - cluster setup 86, 112
 - joining remote clients 101
 - verifying remote client access 102
 - Xgrid administration 26, 27, 34, 54

L

- LDAP (Lightweight Directory Access Protocol)
 - service 86, 98
- libraries, code 61
- local grids 21
- login, SSH 42
- logs 37
- loosely coupled computations 62

M

- Mac OS X
 - agent setup 33, 42
- Mac OS X Server
 - agent setup authentication
 - options 32
 - high performance computing 59, 60
 - software requirements 72
- management computer 78, 81
- memory
 - Xgrid requirements 18
 - Xserve systems 61
- message-passing interface. *See* MPI
- mounting
 - cluster nodes 75
- MPI (message-passing interface) 48

N

- name server 88
 - See also* DNS
- naming conventions, nodes 99, 111
- NAT (Network Address Translation) 90
- NetBoot service 43
- NetInstall 43
- Network File System. *See* NFS
- networks
 - cluster connections 71
 - controller hosting 28
 - grid type 21
 - private 71, 90, 101
 - public 71

- See also* Ethernet
- network services
 - DHCP 87, 98
 - DNS 28, 84, 85
 - NAT 90
 - VPN 90, 101
 - See also* IP addresses
- NFS (Network File System) 90, 93
- nfsd daemon 90
- nodes
 - cluster arrangement 60
 - controller 22
 - firewall settings 88
 - head 21
 - joining to Kerberos realm 112
 - LDAP record 98
 - mounting 75
 - naming 99, 111
 - NAT settings 90
 - overview 18
 - setup 98
 - VPN connection 101

O

- Open Directory 84, 86, 99
- Open Directory master 26, 86

P

- passwords 26, 27, 30, 34
- PdUs (power distribution units) 76
- permissions, administrator 36
- ports
 - Ethernet 77, 81
 - firewall 51, 88
- power considerations 68, 76
- power distribution units. *See* PdUs
- preferences
 - remote setup 42
 - Sharing 33
 - System Preferences 42
- private network 71
 - See also* VPN
- privileges, administrator 36
- problems. *See* troubleshooting
- protocols
 - DHCP 87, 98
 - LDAP 86, 98
- public network 71

R

- RAM (random-access memory) 18
- rated power consumption 68
- realms. *See* Kerberos
- remote server administration 42
 - See also* Apple Remote Desktop

- rendering images 62, 105
- requirements
 - cluster 67, 68, 73
 - hardware 67, 68, 69, 70
 - software 72
 - Xgrid administration 18, 24
- research-related grid projects 18, 19

S

- SACLs (service access control lists) 35
- scp tool 42
- search base, LDAP 86
- secure SHell. *See* SSH
- security
 - administrator permissions 36
 - controllers 88
 - firewall service 28, 51, 88
 - See also* access; authentication
- serial number 72, 75, 99
- Server Admin 99
- serveradmin tool 38
- Server Assistant 81
- servers, remote 42
 - See also* Apple Remote Desktop
- Server Tools 72
- service access control lists. *See* SACLs
- Service Configuration Assistant 28, 29
- setup procedures. *See* configuration; installation
- share points, location of 92
- Sharing preferences 33
- single sign-on (SSO) authentication 26, 27, 34, 54
- software
 - cluster setup 81
 - requirements 72
- space requirements 70
- SSH (secure SHell host) 42, 51, 114
- Static Maps 99
- subnets 24, 87
- supercomputing 17
- System Preferences 42

T

- tail tool 37
- tasks 18, 22, 52, 104
 - See also* jobs
- temperature, operating 69
- troubleshooting
 - agents 51
 - authentication 54

- firewall ports 51
- multi-CPU machines 52
- platform considerations 53
- SSH 51
- typical power consumption 68

U

- uninterruptible power supply. *See* UPS
- UNIX 53, 61
- UPS (uninterruptible power supply) 76
- user accounts, setup 94
- User Datagram Protocol. *See* UDP
- users
 - management of 40
 - volunteer grid projects 18, 19
 - See also* clients; user accounts

V

- ventilation of hardware 77
- Virtual Private Network. *See* VPN
- VPN (Virtual Private Network) 90, 101

X

- Xcode 105
- Xgrid
 - advantages 20
 - application support 53
 - components 22, 23, 24
 - introduction 17, 18, 20, 21
 - management of 37
 - overview 9
 - planning for 26
 - requirements 18, 24
 - setup 25, 30, 91
 - starting 28, 31
 - status checking 37
 - stopping 38
 - See also* agents; clusters; grids, computational; jobs
- Xgrid Admin
 - agents 42, 43, 44
 - grid management 45, 46
 - jobs 44, 45
 - overview 39
 - status indicators 40
 - testing clusters 103, 105, 106
- xgridctl tool 42
- xgrid tool 48, 104
- Xserve 60, 62