Mac OS X Server

Xgrid Administration and High Performance Computing

For Version 10.6 Snow Leopard

Apple

1 Infinite Loop
Cupertino, CA 95014
408–996–1010
www.apple.com

019-1425/2009-05-29

Contents

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 10.6 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 later, or server computers running Mac OS X Server v10.4 or later.

What's New in Xgrid Administration

Xgrid Administration offers major enhancements in several key areas:

- New Xgrid redundancy support
- New Xgrid Admin interface

What's in This Guide

This guide includes the following sections:

- 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 High Performance Computing (HPC) and the applications and tools available for administering HPC.

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

Using Onscreen Help

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

To get the most recent onscreen help for Mac OS X Server:

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

  The onscreen help contains instructions taken from *Advanced Server Administration* and other administration guides.

To see the most recent server help topics:

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

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

Documentation Map

Mac OS X Server has a suite of guides that cover management of individual services. Each service may depend on other services for maximum utility. The documentation map below shows some related guides that you may need in order to fully configure Xgrid service to your specifications. You can get these guides in PDF format from the Mac OS X Server Resources website at www.apple.com/server/macosx/resources/.

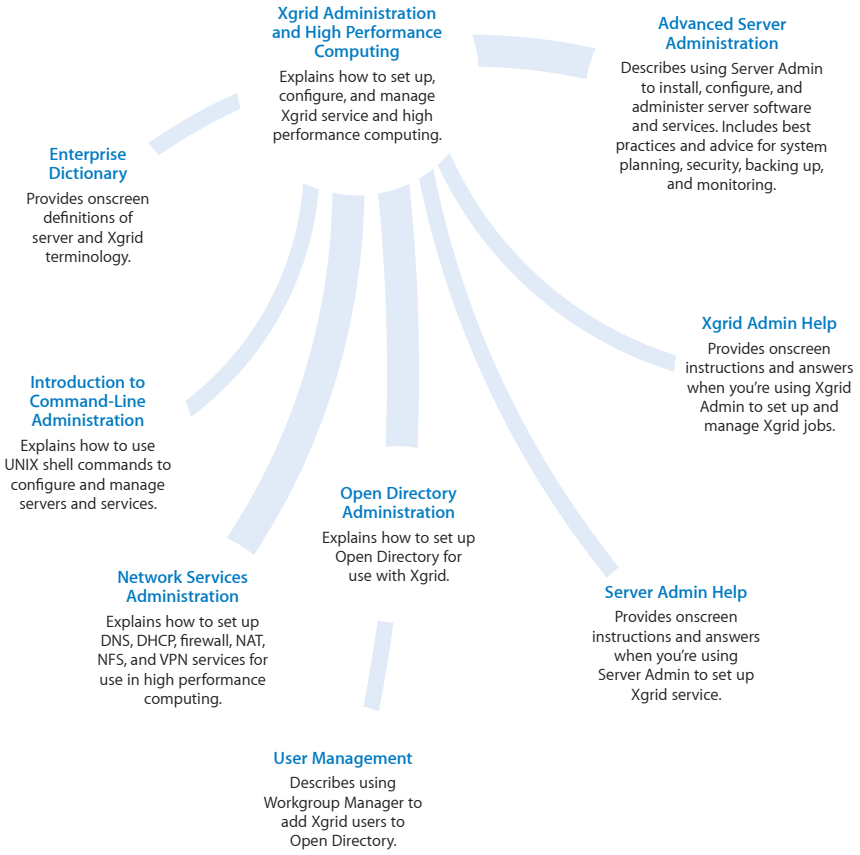

Viewing PDF Guides Onscreen

While reading the PDF version of a guide onscreen:

- Show bookmarks to see the guide's outline, and click a bookmark to jump to the corresponding section.
- Search for a word or phrase to see a list of places where it appears in the guide. Click a listed place to see the page where it occurs.

- Click a cross-reference to jump to the referenced section. Click a web link to visit the website in your browser.

Printing PDF Guides

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

- Save ink or toner by not printing the cover page.
- Save color ink on a color printer by looking in the panes of the Print dialog for an option to print in grays or black and white.
- Reduce the bulk of the printed document and save paper by printing more than one page per sheet of paper. In the Print dialog, change Scale to 115% (155% for *Getting Started*). Then choose Layout from the untitled pop-up menu. If your printer supports two-sided (duplex) printing, select one of the Two-Sided options. Otherwise, choose 2 from the Pages per Sheet pop-up menu, and optionally choose Single Hairline from the Border menu. (If you're using Mac OS X v10.4 or earlier, the Scale setting is in the Page Setup dialog and the Layout settings are in the Print dialog.)

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

Getting Documentation Updates

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

- To view new onscreen help topics for a server application, make sure your server or administrator computer is connected to the Internet and click "Latest help topics" or "Staying current" in the main help page for the application.
- To download the latest guides in PDF format, go to the Mac OS X Server Resources website at www.apple.com/server/macosx/resources/.
- An RSS feed listing the latest updates to Mac OS X Server documentation and onscreen help is available. To view the feed, use an RSS reader application such as Safari or Mail and go to:

  feed://helposx.apple.com/rss/snowleopard/serverdocupdates.xml

Getting Additional Information

For more information, consult these resources:

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

- *Mac OS X Server website* (www.apple.com/server/macosx/)—enter the gateway to extensive product and technology information.
- *Mac OS X Server Support website* (www.apple.com/support/macosxserver/)—access hundreds of articles from Apple's support organization.
- *Apple Discussions website* (discussions.apple.com/)—share questions, knowledge, and advice with other administrators.
- *Apple Mailing Lists website* (www.lists.apple.com/)—subscribe to mailing lists so you can communicate with other administrators using email.
- *Apple Training and Certification website* (www.apple.com/training/)—hone your server administration skills with instructor-led or self-paced training, and differentiate yourself with certification.

Part I: Xgrid Administration

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

Introducing the Xgrid Service

1

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 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 with a minimum of 256 MB of RAM, or Mac OS X Server v10.5 or later with 1 GB 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), or Mac OS X v10.4 or later with 512 MB or RAM.

All Xgrid participants must have a network connection. As always, the more RAM a system has, the better it performs, especially 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 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 19.

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

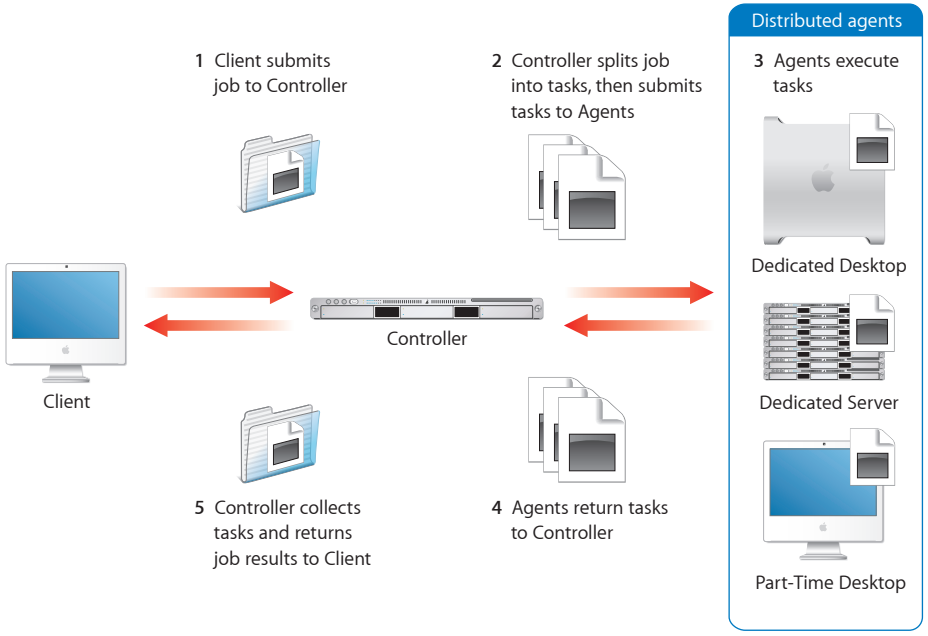

Xgrid has no limitations on the amount of computational power it can support. The performance of the grid depends on the systems participating, the software running, and the network, among other factors. However, individual applications strongly influence the performance of the grid.

You determine if an application is improved by being deployed on a computational grid. In the best case, application performance might 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 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 19
- "Local Grids" on page 20
- "Distributed Grids" on page 20

Xgrid Clusters

Computational clusters are sets of systems dedicated to computation. In a cluster, systems are typically colocated 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 might 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.

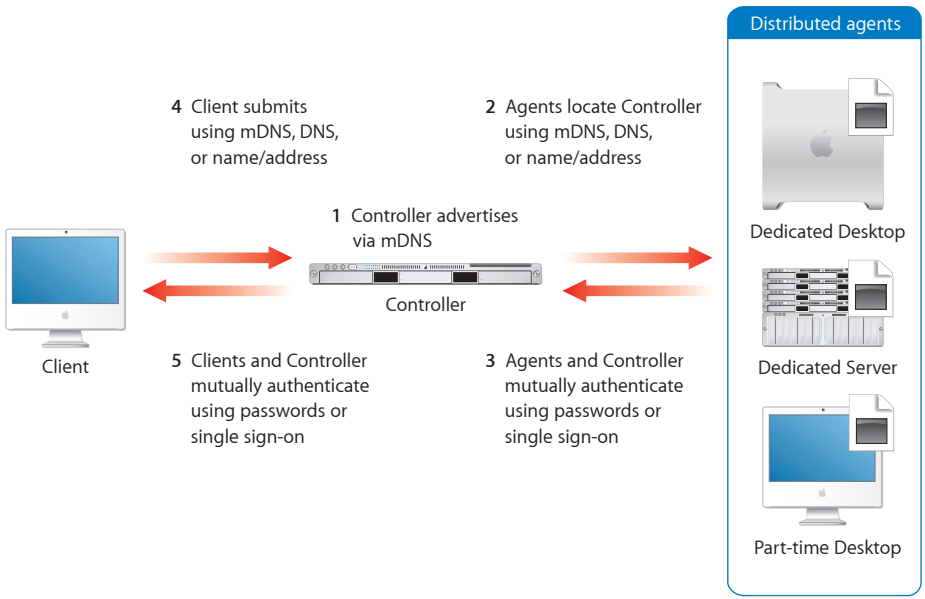

The primary components of a computational grid perform the following functions:

- An agent runs one task at a time per CPU; therefore, a multi-processor computer can run multiple 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 30.

For information about managing agents, see "Managing Xgrid Agents" on page 41.

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 32.

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 29.

For information about managing controllers and grids, see "Managing an 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 can 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.

Setting Up and Configuring the Xgrid Service

2

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 an ad hoc grid or a centrally managed cluster.

Setup Overview

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

Step: 1 Before you begin. See "Before Setting Up Xgrid Service" on page 25. Identify the Xgrid environment you need. Before configuring Xgrid, you must make some decisions about the grid.

Step: 2 Turn Xgrid on. Prior to configuring, turn on Xgrid service. See "Turning Xgrid On" on page 27.

Step: 3 (Optional) Use the Xgrid service configuration assistant to configure Xgrid. You can configure Xgrid using the Xgrid service configuration assistant. This assistant helps with Xgrid configuration by automating many settings. See "Configuring Xgrid with the Xgrid Service Configuration Assistant" on page 27.

Step: 4 Configure Xgrid controller settings. Configure your server as an Xgrid controller using Server Admin. See "Configuring Controller Settings" on page 29.

Step: 5 Start Xgrid. Start Xgrid on the server using Server Admin. See "To see the most recent server help topics:" on page 10.

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 30.

Step: 7 Configure 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 31.

Before Setting Up Xgrid Service

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

- The kind of authentication to use. See "Using Onscreen Help" on page 10.
- Where to host your controller. See "Getting Additional Information" on page 12.
- How you will manage the controller. See "Managing Xgrid" on page 35 and "Monitoring Grid Activity" on page 45.

Authentication Methods for Xgrid

You can configure Xgrid with or without authentication. If you require 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 you enter in Server Admin for the controller must match those you enter 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 32.

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 are 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 might 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 arrange 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 means the controller must be placed outside an organizational firewall (or inside a buffer zone); otherwise, you must open 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 On

Before you can configure Xgrid settings, you must turn Xgrid 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 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 Configuration Assistant

Use the Xgrid 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 configuration assistant:

1. Open Server Admin and connect to the server.
2. Click the triangle at 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 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 the Xgrid Configuration Assistant

Use the Xgrid 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 configuration assistant:

1 Open Server Admin and connect to the server.

2 Click the triangle at 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 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 using your settings.

11 Click Close.

Setting Up Xgrid

You set up Xgrid by configuring two groups of settings on the Settings pane for Xgrid 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 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 "What's in This Guide" on page 9.

Xgrid and Multiple Network Interfaces

On a server with multiple network interfaces, Mac OS X Server makes Xgrid available over all interfaces. You can't configure Xgrid 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 at 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 32.

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 32.

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 32.

Starting Xgrid

Use Server Admin to start Xgrid.

Xgrid 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 30 and "Configuring Controller Settings" on page 29.

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

To start Xgrid:

1 Open Server Admin and connect to the server.

2 Click the triangle at 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 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 at 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 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 32.

   - **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 32.

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 32.

   - **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 32.
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 "Using Onscreen Help" on page 10.

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 at 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 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 30 and "Configuring Controller Settings" on page 29.

To configure password authentication:

1 Open Server Admin and connect to the server.

2 Click the triangle at 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 Xgrid 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:

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:

- Select "Allow only users and groups below."
- Click the Add (+) button to open the Users and Groups window.
- Drag users and groups from the Users and Groups window to the list.

8 Click Save.

Setting Xgrid SACL Permissions for Administrators

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

To set administrator SACL permissions for Xgrid:

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 window by clicking the Add (+) button.
7. From the Users and Groups window, 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

This section describes typical day-to-day tasks you might perform after you set up Xgrid on your server. For information about initial setup, see "Setting Up Xgrid" on page 29.

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

Viewing Xgrid Status

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

To view Xgrid status:

1. Open Server Admin and connect to the server.
2. Click the triangle at 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 system, controller, and agent logs.

Use the View pop-up menu to choose which log to view.

Viewing Xgrid Logs

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

To view logs:

1 Open Server Admin and connect to the server.

2 Click the triangle at 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.

From the command line:

- To view Xgrid log entries:

```
$ tail /var/log/system.log
```

For information about `tail`, see its man page. For the basics of command-line tool usage, see *Introduction to Command-Line Administration*.

Stopping Xgrid

You use Server Admin to stop Xgrid.

To stop Xgrid:

1 Open Server Admin and connect to the server.

2 Click the triangle at 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:

- To stop Xgrid:

```
$ sudo serveradmin stop xgrid
```

For information about `serveradmin`, see its man page. For the basics of command-line tool usage, see *Introduction to Command-Line Administration*.

About Xgrid Redundancy

Mac OS X Server v10.6 supports Xgrid redundancy, which lets you configure an Xgrid controller with backup Xgrid controllers. This type of redundancy is referred to as active-passive redundancy.

With Xgrid redundancy you can have one active controller, which is the primary Xgrid server that services all Xgrid job requests. The active controller can have an undefined number of passive controllers that are promoted to active and begin to service requests if something happens to the active Xgrid controller.

A passive controller can be instructed to promote itself to active by an authorized user or service. The promotion instruction can be sent manually to a server's controller using Server Admin. Authorized remote monitors can also send promotion instructions. An example of an initial external remote monitor would be Podcast Producer. The passive controllers wait indefinitely to be promoted.

Remote monitors observe the status of the active and passive controllers and, based on their own policies, decide when a promotion or demotion instruction should be sent and the controller to send it to. If there are no active controllers, no jobs are processed.

When a passive controller receives promotion instructions the passive controller updates to the database to designate itself as an active controller. At this point there are two designated active controllers: the original active controller and the promoted active controller.

The next time the original active controller reads the database it discovers that it is no longer the only active controller and, based on time stamps, determines that it should demote itself. The original active controller updates the database to designate itself as passive, drops connections with clients and agents, and then restarts.

The active controller maintains a list of available passive controllers. The passive controller updates a time stamp in their database periodically so that the active controller knows they are still online.

Agents and clients can subscribe to the list of passive controllers. When an agent loses its connection with the active controller, it asks each passive controller if it is now the active controller. The agent keeps doing this until it finds an active controller.

Based on the unique identifiers and how many logical active controllers are found, the agent can choose which active controller to use at that point.

By using Xgrid redundancy you can improve server reliability and lessen maintenance down-time. Your Xgrid service will become more reliable because you have additional servers ready to process jobs if a server fails or is shut down. This redundancy prevents administrators from canceling jobs or disconnecting users from a controller while doing maintenance on the Xgrid controller.

Setting Up Xgrid Redundancy

When setting up Xgrid redundancy there is only one active controller but there can be any number of passive controllers. To provide redundancy, all controllers must share a database and have access to the same location in a shared file system. Also, clients and controllers (active and passive) must use Kerberos authentication.

To set up Xgrid redundancy, complete these steps.

Step: 1 Setup an Active Controller To set up an active controller, you configure it to store its data in a shared file system. When the controller starts, it creates a database in the shared file system location and designates itself as the active controller.

Step: 2 Setup a Passive Controller To set up a passive controller, you configure it to store its data in the same location as an active controller. When the controller starts, it opens the database in the shared file system location and designates itself as a passive controller. Each controller has a unique identifier separate from its address or hostname.

Managing a Grid Using Xgrid Admin

3

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 v10.5 or later.

Xgrid Admin Overview

Xgrid Admin is a tool you use to monitor grids and manage agents and jobs. You can add controllers and agents to monitor and specify agents that have not joined a grid. You also use Xgrid Admin to pause, stop, or restart jobs.

You can manage 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.

For more information, see "Using Xgrid Admin" on page 39.

Using Xgrid Admin

Xgrid Admin enables you to monitor grids and manage agents and jobs. You can:

- Check the status of a grid and its activity, including the number of agents working and available, the 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.

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," and "Submitting a Job" on page 48.

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 an Xgrid Controller

In general, you manage an 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.

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, click the Action pop-up menu, and select Connect.
2. If necessary, select your 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 the Action pop-up menu and select 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 the Add (+) button and select 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 the Action pop-up menu and select Remove Controller.

Managing Xgrid 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/technology/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, 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 at 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 are 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 at 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 Xgrid Jobs

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

Note: 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 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.

Managing Grids

You use Xgrid Admin to manage a grid.

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 and select Add Grid.
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.

Planning and Submitting Xgrid Jobs

4

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 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 might 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 is in one single large job, with numerous small tasks.
- Everything is 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.)
- An entire workflow is 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 a job 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 do 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/.

When you submit a job to a grid make sure you use a universal binary. This assures that your job has the correct architecture no matter what architecture the grid agents provide.

Also, make sure you set your deployment target correctly. For example, if you are building a tool for Mac OS X v10.3 you must build with Mac OS X v10.3 as your deployment target.

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 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 Xgrid 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>
```

Solving Xgrid Problems

5

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 specific 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 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 relevant 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 updated for Mac OS X v10.5 or later), at http://www.novajo.ca/xgridagent/

If the Xgrid Controller Must Be Restarted

When the Xgrid controller is restarted by Server Admin, the `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) 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 Kerberos domain controller 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 Kerberos domain controller 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 Kerberos domain controller 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 provides news, posted questions and threads, and the ability to post 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

II

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

Introducing High Performance Computing (HPC)

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 v10.6 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 v10.6 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 multiprocessor and multicore 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 using difficult programming techniques or expensive development tools.

Mac OS X Server also includes Xgrid, an integrated distributed resource manager for 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: a cluster controller and one or more compute nodes, as shown in the following illustration:

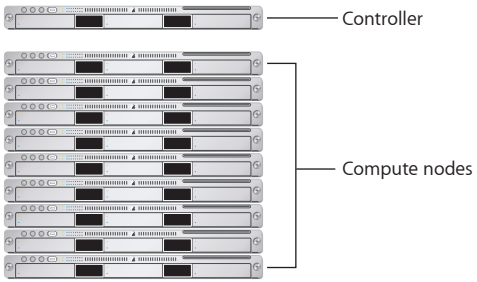

Xserve cluster

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 are 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 an Open Brand UNIX 03 Registered Product, conforming to SUSv3 and POSIX 1003.1 specifications for the C API, shell utilities, and threads. It can compile and run all 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 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.

  The Apple Workgroup Cluster is a configured 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 might benefit from another interconnect technology.

Setting Up an HPC Cluster

7

Use this chapter to learn about the process of setting up an HPC cluster.

You 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 consists of a controller and a number of compute nodes.

The compute nodes are connected to the controller via a private (isolated) Ethernet network switch. The controller is connected to the private Ethernet network and a public network, potentially the Internet. The controller also provides a shared file system to compute nodes.

The controller provides a number of services to the compute nodes:

- A firewall isolates the controller and compute nodes from the public network, protecting against unwanted access. Access to the private network from outside the firewall requires 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 allows the compute nodes to communicate with each other and external networks.
- Open Directory contains user account information, including usernames and passwords, and makes these accounts available to compute nodes. Using Kerberos with Open Directory provides single sign-on capability, reducing the number of times a user must enter passwords to access cluster resources.
- Open Directory also publishes 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, is mounted on each compute node.
- The controller hosts the Xgrid controller service.

Cluster Setup Overview

Here is a summary of what you do 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 HPC 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 HPC Cluster Controller."

Use Server Admin to configure and start the following services:

- DNS service. See "Configuring the DNS Service" on page 76.
- Open Directory service. See "Configuring the Web Service" on page 84.
- DHCP service. See "Configuring the DHCP Service" on page 80.
- Firewall service. See "Configuring Firewall Settings on the Cluster Controller" on page 81.
- NAT service. See "Configuring NAT Settings on the Cluster Controller" on page 83.
- NFS service. See "Configuring NFS" on page 83.
- VPN service. See "Configuring VPN" on page 83.
- Xgrid service. See "Configuring Xgrid" on page 84.

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 85.

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 86.

Step 6: Create network user accounts. Use Workgroup Manager to create network user accounts for cluster users. See "Creating User Accounts" on page 87.

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 configure themselves when they start up for the first time.

See "Creating an Auto Server Setup Record for Compute Nodes" on page 89 and "Verifying LDAP Record Creation" on page 90.

Step 8: Set up compute nodes. Start compute nodes to begin the Auto Server Setup process. They'll configure themselves and then restart. See "Setting Up Compute Nodes" on page 91.

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. See "Configuring Cluster Nodes" on page 92 and "Joining a Remote Client to the Kerberos Realm" on page 94.

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 an HPC Cluster."

Identifying Prerequisites and System Requirements

8

Use this chapter to learn the prerequisites and requirements for setting up an HPC cluster and to 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 v10.6 or later. If not, you must install a clean version of Mac OS X Server v10.6 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 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 include the following:

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 might not be high enough, depending on your configuration. For example, if your cluster uses Xserve RAID systems, or third-party hardware, you must include their power consumption requirements.

Cooling Requirements

It's very important to keep Xserve computers running at normal operating temperatures (see www.apple.com/xserve/specs.html). If servers overheat they will shut down and 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, especially 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 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 of 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:

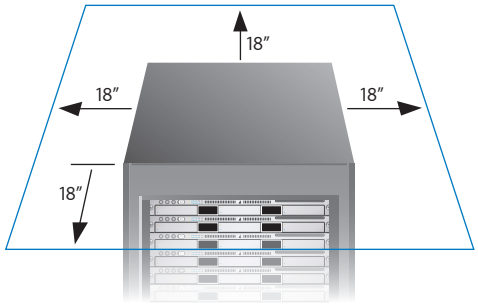

You should have enough space to open the rack's door, slide out systems, and perform other routine maintenance tasks.

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 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.

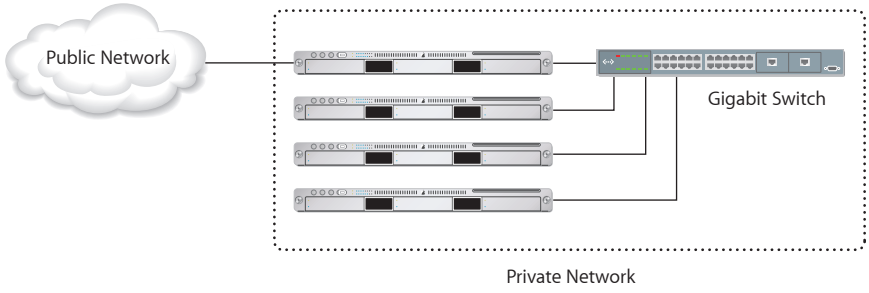

Software Requirements

You need:

- A site-licensed copy of Mac OS X Server v10.6 or later.
- 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, 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-xxx-x, where x is a letter and 9 is a digit.

Apple Remote Desktop

Configuration and administration of your cluster is 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 to rapidly install software.

Server Tools

If you use 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:

- Reserve a range of IP addresses for the private network. A number of nonroutable IP address ranges are available. These addresses cannot be used with the Internet without Network Address Translation (NAT), which is 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 in the home, and because your users might want to connect to your cluster from behind one of these devices, choose a range less likely to exist on 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 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, 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 address, 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 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 v10.6 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
serial_number_0	Cluster controller
serial_number_1	Compute node 1
serial_number_2	Compute node 2
. . .	. . .

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

- The unit's back panel

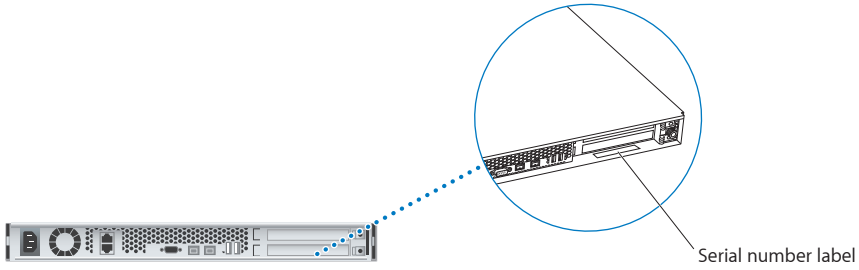

- The unit's interior.

  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.

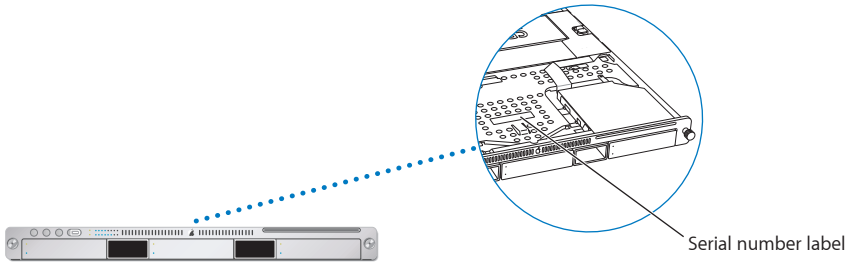

- The large pull-out plastic tab on Xserve computers with Intel processors
- The cardboard shipping box

  You can use a barcode scanner on the box label to get the serial number.

3 Using 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 below) by connecting port 1 on the cluster controller to the public network and port 2 to the private network.

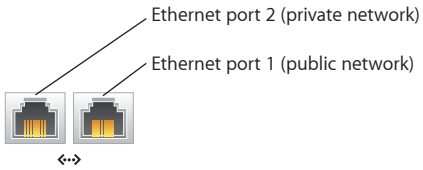

5 Connect Ethernet port 1 on the remaining nodes in the cluster to the private network, in order.

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.

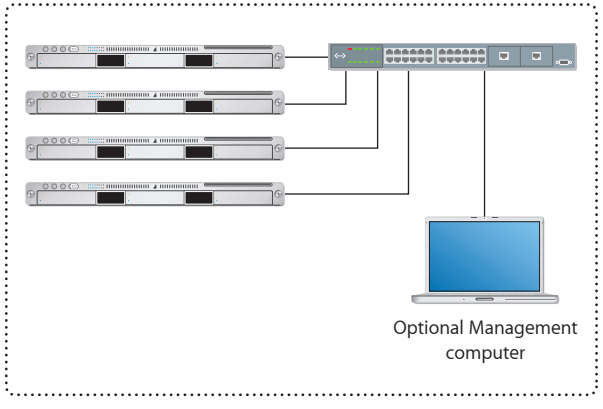

2 Start the management computer.

3 Disable AirPort and any network connection other than the one you're 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.

Install the Mac OS X Server tools and applications 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, 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, 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 used in this guide (172.16.1.1 – 172.16.1.199 for compute nodes, 172.16.1.254 for the controllers), configure your management computer to use 172.16.1.253.

After you connect to the private network, install the server administration tools mentioned in this guide (Server Assistant, Server Admin, Workgroup Manager, and Xgrid Admin) and use them 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 or later, to control the nodes via the network, using the server administration tools directly on the remote nodes.

Setting Up the HPC Cluster Controller

10

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 *Advanced Server Administration*.

   If you are using 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 Keyboard screen, select the keyboard layout for the server and click Continue.
4. In the Serial Number screen, enter a volume license Mac OS X Server serial number and click Continue.
5. In the Transfer an Existing Server screen, select "Set up a new server" and click Continue.
6. In the Registration Information screen, fill out the form or press Command–Q and click Skip.

7. In the Time Zone screen, select your time zone form the Closest City pop-up menu and click Continue.
8. In the Administrator Account screen, create the user account to administer the cluster controller with (for example, Administrator) and click Continue.
9. In the Xsan screen, select "Don't configure Xscan now" and click Continue.
10. In the Network screen, configure Ethernet 1:
    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.
11. In the Network screen, configure Ethernet 2:
    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, 172.16.1.1).
    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 Continue.
12. When the alert message "The router address is not valid" appears, click Ignore.
13. When the alert message "The host portion of the router is all zeros" appears, click Ignore.
14. In the AirPort Management screen, disable Airport Management and click Continue.

15 In the Network Names screen, click Edit to set the Primary DNS name of the controller and click Allow Editing.

Set the Primary DNS name (Public name) to controller.cluster and set the Computer Name (Private name) to controller then click Continue.

16 In the Users and Groups screen, select Configure Manually then click Continue.

17 In the Connect to a Directory Server screen, don't connect to a directory server click Continue.

18 In the Directory Services screen, select the No checkbox next to "Set up a Open Directory master" then click Continue.

19 In the Review screen, review and confirm your selections then click Set Up.

20 In the Setting Up screen, click Go.

21 When your server starts, click Apple menu > Software Update and install updates.

Some updates may require you to reboot your server to complete the update installation. Continue installing updates until there are no updates available for your server.

22 Launch Server Admin, and then (if prompted) enter the administrator user name and password.

23 Click Connect then click Choose Configured Services.

24 Select the checkboxes to enable the following services: DHCP, DNS, Firewall, NAT, NFS, NetBoot, Open Directory, VPN, Web, and Xgrid.

25 Click Save.

26 To reveal the enabled services, expand the triangle next to the controller.local file in the Servers list.

Configuring the 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 open.

2 If necessary, click the triangle at 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 Delete all zones by clicking the Delete (–) button.

9 Click the Add Zone button, then select "Add Primary Zone (Master)."

A default zone named example.com is created.

10 Select the default example.com zone.

11 Change the primary zone name to your private DNS domain.

The primary zone name must end with a period (for example, "cluster.").

12 Set Admin Email to the mail address of the person who should be notified of DNS errors (for example, administrator@controller.cluster).

13 Click the Add (+) button next to the Nameservers list.

14 Enter the private DNS host name as the zone name, ending the zone name with a period (for example, cluster).

15 Enter the primary DNS name as the Nameserver Hostnamem, ending the Nameserver Hostname with a period (for example, controller.cluster).

16 Click Save.

17 Select the cluster DNS zone.

18 Click the triangle at the left of the cluster DNS zone.

19 Click Add Record, then select "Add Machine (A)."

20 Change the Machine Name field to the private hostname of the controller (for example, controller).

21 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).

22 Click Save.

23 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 10.0.2.2, node2 assigned to 10.0.2.3, and so on.

24 Click Settings.

25 Click the Add (+) button (below the Forward IP Addresses list).

26 Enter an forwarding IP Address (for example, 172.16.1.255).

27 Click the Start DNS button (below the Servers list).

The DNS service status indicator turns green when the service starts.

28 From the Apple Menu (), open System Preferences (/Applications/System Preferences).

29 Click Network.

30 Select the Ethernet 1 interface.

31 In the DNS Server field, enter the public IP address of the controller (for example, 10.0.2.199).

32 In the Search Domains field, enter the private DNS domain (for example, cluster).

33 Click Apply.

34 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. 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 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 the Cluster Controller as an Open Directory Master

Use Server Admin to configure Open Directory on the cluster controller. Open Directory is responsible for authenticating users, publishing server setup configurations, and publishing network share automount records.

To configure Open Directory settings:

1. Open Server Admin if it is not 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 Assistant.
4. Select Set up an 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 set to be 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 set 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 populated, your DNS settings might not be configured properly. If so, click the Cancel button and redo the steps listed in "Configuring the DNS Service" on page 76.
7. Confirm the settings, then click Continue.
8. When the service configuration assistant completes, click Done.
9. Verify that 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 created. You can also use the Console (located in /Applications/Utilities/) or Terminal with the command "tail -f/Library/Logs/slapconfig.log." You can ignore warnings such as the following in the log:

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

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

Configuring the DHCP Service

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

To configure the DHCP service:

1. Open Server Admin if it is not open.
2. In the controller's list of services, click DHCP.
3. Click Subnets.
4. Remove all subnets.
5. Click the Add (+) button to create a new subnet for Ethernet 2.
6. Click General.
7. In the Subnet Name field, enter a subnet name (for example, Cluster Private Network).
8. In the Starting IP Address field, enter the first IP address in the private network range available for compute nodes (for example, 172.16.1.2).
9. In the Ending IP Address field, enter the last IP address in the private network range available for compute nodes (for example, 172.16.1.253).

   Note: Leave some addresses unused at the end of the range for other devices and VPN connections.
10. In the Subnet Mask field, enter the subnet mask for your private network (for example, 255.255.255.0).
11. From the Network Interface pop-up menu, select en1 if it is not selected.

    This menu shows the UNIX name for the port. The UNIX name for Ethernet 2 should be en1.
12. In the Router field, enter the private IP address of the cluster controller (for example, 172.16.1.1).
13. Set the lease time for the IP addresses served by the DHCP service to at least 1 month.
14. Click Save.
15. Click DNS below the Subnets list.
16. In the DNS Servers field, enter the public address of the cluster controller (for example, 10.0.2.199).
17. In the Default Search Domain field, enter the DNS domain for your private network (for example, cluster).
18. Click Save.
19. Click LDAP.
20. In the Server Name field, enter the fully qualified DNS name of the cluster controller (for example, controller.cluster).

21 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.

22 Verify the Server Name and Search Base fields.

Errors in the LDAP configuration of 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 search base settings.

23 Select the Enable checkbox to the left of the subnet you just created.

24 Click Save.

25 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 "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, delete the first entry and enter your public 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).

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, delete the first entry and enter your private 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).

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 to these ports."

17 Select the following ports (in addition to what's already selected):

- VPN PPTP—Point-to-Point Tunneling Protocol (1723)
- VPN PPTP: GRE—Generic Routing Encapsulating protocol

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."

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."

24 Click Advanced.

25 Click the Add (+) button (below the Advanced Rules list).

26 From the Action pop-up menu choose Allow.

27 From the Protocol pop-up menu choose UDP.

28 From the Source Address pop-up menu choose "any."

29 From the Destination Address pop-up menu choose "any."

30 From the Destination Interface pop-up menu choose Other.

31 Click OK.

32 Click Save.

33 Click the Start Firewall button (below the Servers list).

Configuring NAT Settings on the Cluster Controller

Network Address Translation (NAT) allows computer 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) are 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).

6 If an alert message appears explaining that the service port may be restricted, click OK.

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).

6 If an alert message appears explaining that the service port may be restricted, click OK.

Configuring VPN

Configure VPN 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, 10.0.2.200).

5 In the Ending IP address field, enter the last private IP address you want to assign to remote VPN clients (for example, 10.0.2.229).

6 Click Save.

7 Click the Start VPN button (below the Servers list).

8 If an alert message appears explaining that the service port may be restricted, click OK.

Configuring the Web Service

Configure the web service to enable web based monitoring of cluster.

To configure the web service:

1 In the controller's list of services, click Web.

2 Click the Start Web button (below the Servers list).

3 If an alert message appears explaining that the service port may be restricted, click OK.

Configuring NetBoot

Configure NetBoot to manage NetBoot install images for your Xgrid nodes. The NetBoot install images are used to install Mac OS X images on nodes.

To configure NetBoot:

1 In the controller's list of services, click NetBoot.

2 Click Settings then click General.

3 Select the Ethernet 2 checkbox.

4 From the Volumes list, select the Images and Client Data checkboxes for the volumes you will use for storing your NetBoot install images.

5 Click Save.

Configuring Xgrid

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

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 Xgrid:

1. In the controller's list of services, click Xgrid.
2. Click Overview.
3. Click Configure Xgrid Service.

   The service configuration assistant launches.
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 Open Directory.
7. Click Continue.
8. Verify that the Xgrid settings include the correct Kerberos realm (for example, CONTROLLER CLUSTER).
9. Click Continue.
10. After Xgrid 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 you prepare a data drive you should protect your data by using a mirrored RAID set, referred to as RAID 1. You can use the Disk Utility application to create the 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 Disk Utility (in /Applications/Utilities).
2. From the drive list on the left, click a drive to use in the RAID set.
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 list at the center of the pane.
6. For each disk you dragged to the 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 "Mac OS Extended (Journaled)" or "Mac OS Extended (Case-sensitive, Journaled)."

If you 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 set 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 Disk Utility.

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 and group 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, 10.0.2.0).

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 86), restart it now.

2 Log in using your administrator account.

3 Open Workgroup Manager (located in /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.

Setting Up Compute Nodes

11

Use this chapter to 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 *Advanced 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

You can use the Auto Server Setup profile to stream line the installation of your node on your Xgrid cluster.

To configure an auto server setup profile:

1. On the cluster controller, open Server Admin and click Server menu.
2. Choose Create Auto Server Setup Profile.

   Server Assistant opens.
3. In the Keyboard screen, select your keyboard layout then click Continue.
4. In the Serial Number screen, enter a site-licensed Mac OS X Server serial number and click Continue.

   If you don't have a site-licensed number you must manually enter unique serial numbers for each compute node after it is configured.
5. In the Time Zone screen, choose your time zone form the Closest City pop-up menu and click Continue.
6. In the Administrator Account screen, create the account you'll use to administer the compute nodes (for example, Administrator) and click Continue.

7. In the Network screen, click the Add (+) button.
8. From the Type pop-up menu, choose Ethernet
9. In the Service Name field, enter Ethernet.
10. In the Interface field, enter en0.
11. Leave the Hardware Address field blank.
12. Click Create and then click Continue.
13. In the Network Names screen, leave the Primary DNS Name and Computer Name fields blank.
14. Click Continue

    If a warning appears indicating that you left some fields blank, click Skip.
15. In the Users and Groups screen, select Configure Manually, then click Continue.
16. In the Services screen, deselect all services checkboxes, then click Continue.
17. In the Save Configuration screen, choose None from the Encryption pop-up menu.
18. Select the Restrict for use with certain computers checkbox, then click Edit.
19. Change the statement to read "IP address begins with 10.0.2," and then click OK.
20. Click Save.
21. From the Where pop-up menu choose Desktop then click Save.
22. Click Start Over, and then 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=autoserversetup,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.

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.

Each node in a cluster should 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 configures itself and reboots.

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.

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.
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. To verify the serial number, click General.
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 104.
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.
18. Set the realm to your Kerberos realm (for example, CONTROLLER.CLUSTER).
19. Enter the Open Directory administrator user name and password.
20. Click Refresh below the Servers list.

If the node has joined the Kerberos realm, the Join Kerberos button and associated text disappears.

21 In the Servers list select the node's Xgrid service.

22 Click Overview.

23 Click Configure Xgrid Service.

The Xgrid Service Configuration Assistant appears.

24 Click Continue, then select "Join a grid."

25 Click Continue.

26 In the "Use controller with hostname" field, enter the controller's private DNS name (for example, controller.cluster).

27 Click Continue.

28 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).

29 Click Continue.

30 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 or Mac OS X v10.6 clients. For other operating systems, or older versions of Mac OS X, consult the relevant 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 87.

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 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.

Creating and Distributing a NetInstall image

Use System Image Utility to create a NetInstall image that you can use to install software on nodes over the network. You can find this application in the /Applications/Server/ folder.

When creating an workflow for your image the Define Image Source action must be at the beginning of the workflow and the Create Image action must be at the end of a workflow.

To create a NetInstall image:

1. Log in as an administrator user.
2. Open System Image Utility (in the /Applications/Server/ folder).
3. From the left sidebar, select the image source.

   If no image sources are listed, make sure you inserted a valid Mac OS X v10.6 or later installation DVD or mounted a valid Mac OS X v10.6 or later boot volume.

   Note: To create an image, you must have valid Mac OS X v10.6 image sources (volumes or installation DVDs). You cannot create an image of the startup disk you are running on.
4. Select NetInstall Image and click Customize.

   The Define Image Source and the Create Image actions are already in the workflow.

5. When the Software License Agreement appears, click Agree.
6. In the Define Image Source action, choose Mac OS X Server Install Disc.

   This action must be at the beginning of all image creation workflows.
7. From the Automator Library, drag the Enable Automated Installation action below the Define Image Source action.
8. Select Named and enter a name for the volume (for example, Server HD).
9. Create a user account using the Name, Short Name and Password fields (for example, Cluster Administrator, clusteradmin, password) and select the "Allow user to administer the computer" checkbox.
10. From the Automator Library, drag the Add User Account action below the Enabled Automated Installation action and enter the following information:
    - **Name** Enter a user name (for example, Cluster Administrator).
    - **Short Name** Enter a short name for the user (for example, clusteradmin).
    - **Password** Enter a password for the user.
11. Select the Allow user to administer the computer checkbox.
12. In the Create Image action, select NetInstall as the Type and choose NetBootSP0 from the Save To pop-up menu.
13. Click Run.
14. Click Save and authenticate if prompted.
15. Quit System Image Utility.
16. Open Server Admin and connect to the controller.
17. In the controller's list of services, click NetBoot.
18. Click Settings then click Images.
19. From the list of images, select the NetInstall image you want as your default image.
20. Select the Enable checkbox for your default image.
21. Click Save.
22. Click the Start NetBoot button (below the Servers list).

Testing an HPC Cluster

12

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

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 a Cluster Using Xgrid Admin

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

To use Xgrid Admin to check a cluster:

1. From the management computer, a VPN client, or the controller, open Xgrid Admin (located in /Applications/Server/).
2. Click the Add (+) button and select Add Controller.
3. From the pop-up menu, choose the controller, then click the Action pop-up menu and select Connect.
4. In the authentication sheet, complete the following:
   - Select "Use Single Sign On Credentials."
   - Click OK.
   - If prompted, enter a cluster account username, the Kerberos realm (for example, CONTROLLER.CLUSTER), and password.
   - 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 and unavailable processors, working processors, 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 50.

9 Monitor the progress of Xgrid jobs by clicking Jobs.

10 Quit Xgrid Admin.

Testing an Xgrid Cluster

To test your cluster, use GridSample, a sample Cocoa application that comes with Developer Tools for Mac OS X v10.6, to submit Xgrid tasks to the controller. This application provides an easy-to-use GUI for Xgrid. On any system that has Mac OS X developer tools installed, the example code for the application is in /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 *Introduction to Command-Line Administration*>.

To test a 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 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 an Xgrid Configuration

Verify that Xgrid is configured and works.

To verify an Xgrid service:

1 Install and configure Xcode developer tools.

Xcode is included with the Mac OS X Server Installation disc. You can download the latest version of Xcode 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 an SSH Connection

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

To verify an 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).

3 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 compute nodes using their private DNS name (for example, `ssh tomclark@node1.cluster` or `ssh tomclark@node1`).

Cluster Setup Checklist

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	
	Power source meets minimum requirements	"Power Requirements" on page 64
	Cooling system meets minimum requirements	"Cooling Requirements" on page 65
	Facility housing the cluster meets minimum weight requirements	"Weight Requirements" on page 66
	Space around the cluster meets minimum requirements	"Space Requirements" on page 66
	Network switches support Gigabit Ethernet and have enough ports	"Network Access Requirements" on page 67
	Mount cluster nodes on the rack	"Network Access Requirements" on page 67
	Connect cluster nodes to a power source	"Preparing Cluster Nodes for Software Configuration" on page 70
	Connect cluster nodes to the private network	"Preparing Cluster Nodes for Software Configuration" on page 70

	For information about this step	Go to
	Software Setup	
	Obtain a static IP address and related network and DNS information	"Network Access Requirements" on page 67
	Obtain a site-licensed serial number	"Volume-Licensed Serial Number" on page 68
	Obtain a copy of Apple Remote Desktop	"Apple Remote Desktop" on page 68
	Record the serial numbers of cluster nodes	"Preparing Cluster Nodes for Software Configuration" on page 70
	Management Computer Setup (Optional)	
	Disable AirPort and other public network connections	"(Optional) Setting Up the Management Computer" on page 72
	Install the latest version of Mac OS X Server tools	"(Optional) Setting Up the Management Computer" on page 72
	Install Apple Remote Desktop	"(Optional) Setting Up the Management Computer" on page 72
	Controller Setup	
	Connect the controller to the public and private network	"Setting Up Server Software on the Cluster Controller" on page 74
	Run Server Assistant and configure public network settings	"Setting Up Server Software on the Cluster Controller" on page 74
	Configure DNS service	"Configuring the DNS Service" on page 76
	Configure Open Directory service	"Configuring the Cluster Controller as an Open Directory Master" on page 79
	Configure DHCP service	"Configuring the DHCP Service" on page 80
	Configure Firewall service	"Configuring Firewall Settings on the Cluster Controller" on page 81

	For information about this step	Go to
	Configure NAT service	"Configuring NAT Settings on the Cluster Controller" on page 83
	Configure NFS service	"Configuring NFS" on page 83
	Configure VPN service	"Configuring VPN" on page 83
	Configure Xgrid service	"Configuring Xgrid" on page 84
	Prepare data drive	"Preparing the Data Drive as a Mirrored RAID Set" on page 85
	Create home directory	"Creating a Home Directory Automount Share Point" on page 86
	Create user accounts	"Creating User Accounts" on page 87
	Compute Node Setup	
	Create auto server setup records	"Creating an Auto Server Setup Record for Compute Nodes" on page 89
	Set up compute nodes	"UNRESOLVABLE CROSS-REFERENCE" on page ###
	Configure cluster nodes	"Configuring Cluster Nodes" on page 92
	Create and verify VPN connection	"Creating and Verifying a VPN Connection" on page 93
	Cluster Testing	
	Check the cluster using Xgrid Admin	"Checking a Cluster Using Xgrid Admin" on page 97
	Test Xgrid cluster	"Testing an Xgrid Cluster" on page 98
	Verify Xgrid configuration	"Verifying an Xgrid Configuration" on page 99
	Verify your SSH connection	"Verifying an SSH Connection" on page 99

Automating Compute Node Configuration

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

For large clusters, you can complete some tasks in this guide 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 the 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.

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, an entry in the results window appears.

8 Close the Send UNIX Command results window.

Configuring Xgrid Agent Settings Using Apple Remote Desktop

To send commands to compute nodes to configure 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 compute nodes.

4 Select 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:ReadyForControllerRoleBasedS
    etup = 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 a 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.

Index

D

E

F

G

H

I