

Xsan 2

Administrator's Guide

🍏 Apple Inc.
© 2008 Apple Inc. All rights reserved.

Under the copyright laws, this manual may not be copied, in whole or in part, without the written consent of Apple.

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

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

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

Apple, the Apple logo, Mac, Mac OS, Macintosh, the Mac logo, Tiger, Xsan, and Xserve are trademarks of Apple Inc., registered in the U.S. and other countries.

Apple Remote Desktop, Finder, Leopard, and Spotlight, are trademarks of Apple Inc.

ADIC and StorNext are registered trademarks of Advanced Digital Information Corporation.

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

PowerPC and the PowerPC logo are trademarks of International Business Machines Corporation, used under license therefrom.

UNIX is a registered trademark of The Open Group.

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

The product described in this manual incorporates copyright protection technology that is protected by method claims of certain U.S. patents and other intellectual property rights owned by Macrovision Corporation and other rights owners. Use of this copyright protection technology must be authorized by Macrovision Corporation and is intended for home and other limited viewing uses only unless otherwise authorized by Macrovision Corporation. Reverse engineering or disassembly is prohibited.

Apparatus Claims of U.S. Patent Nos. 4,631,603, 4,577,216, 4,819,098 and 4,907,093 licensed for limited viewing uses only.

Simultaneously published in the United States and Canada.

019-0964-A/2008-01-16

Contents

Preface	9 About This Book
	10 What's New in Xsan 2
	10 Version Compatibility
	10 Upgrading from an Earlier Version of Xsan
	11 Getting Additional Information
	11 Notation Conventions
Chapter 1	13 Overview of Xsan
	13 Xsan Storage Area Networks
	14 Shared SAN Volumes
	15 Metadata Controllers
	15 Clients
	15 Network Connections
	16 How Xsan Storage Is Organized
	16 LUNs
	17 Storage Pools
	17 Affinities
	18 Volumes
	19 Folders with Affinities
	19 How Xsan Uses Available Storage
	19 Metadata and Journal Data
	19 Striping at a Higher Level
	20 Security
	20 Expanding Storage
	21 Xsan Capacities
Chapter 2	23 Planning a Storage Area Network
	23 Hardware and Software Requirements
	23 Supported Computers
	24 Supported Storage Devices
	24 Fibre Channel Fabric
	25 Ethernet TCP/IP Network
	26 Directory Services

27	Outgoing Mail Service
27	Planning Your SAN
28	Planning Considerations and Guidelines
34	Planning the Ethernet TCP/IP Network
35	Using a Private Metadata Network
35	Using Switches Instead of Hubs
35	Planning the Fibre Channel Network
35	Verifying Base Fibre Channel Performance
36	If Your Fibre Channel Fabric Is Running Slower Than Expected
36	Configuring RAID Systems
36	Installing the Latest Firmware
37	Connecting RAID Systems to an Ethernet Network
37	Choosing RAID Levels for LUNs
37	Adjusting RAID System Performance Settings

Chapter 3

39	Setting Up a Storage Area Network
39	Connecting Computers and Storage Devices
39	Preparing LUNs
40	Using Server Assistant to Configure Controllers
40	Managing Users and Groups with Xsan Admin
40	Using an Existing Open Directory Server
40	Using Another Directory Server
41	Using Xsan Admin
41	Installing Just the Xsan Admin Application
41	Connecting Through a Firewall
41	Xsan Admin Preferences
41	Getting Help
41	SAN and Volume Setup Summary
42	Setting Up an Xsan Volume on a Storage Area Network
51	Setting Up an Xsan Administrator Computer
51	Renaming a SAN
52	Removing a SAN
52	Managing Multiple SANs
52	Setting Up Additional SANs

Chapter 4

53	Managing SAN Storage
53	Adding Storage
54	Preparing LUNs
54	Finding the Drive Modules That Belong to a LUN
55	Adding LUNs to a Storage Pool
56	Rearranging Fibre Channel Connections
57	Adding a Storage Pool to a Volume
58	Adding a Volume to a SAN

59	Setting Up a Folder Affinity
60	Changing a Folder's Storage Pool Affinity
60	Removing an Affinity
61	Changing Advanced Volume Settings
62	Setting the Block Allocation Size
63	Changing the Volume Allocation Strategy
63	Enabling or Disabling Spotlight on a Volume
63	Enabling or Disabling Access Control Lists
64	Changing the Windows ID Mapping
64	Changing Advanced Allocation and Cache Settings
65	Renaming a Volume
65	Changing Storage Pool Settings
65	Choosing the Data Type for an Affinity Tag
66	Setting the Storage Pool Stripe Breadth
67	Checking Volume Fragmentation
67	Defragmenting a Volume
68	Checking the Integrity of a Volume
68	Repairing a Volume
69	Destroying a Volume

Chapter 5

71	Managing Clients and Users
72	Adding a Client
73	Adding an Xsan Serial Number
74	Moving a Client to a Different SAN
75	Mounting a Volume on a Client
76	Changing Mount Options
77	Managing Users and Groups
78	Adding SAN Users
79	Deleting SAN Users
79	Creating Groups
79	Deleting Groups
80	Changing Group Membership
80	Controlling Client and User Access
80	Controlling File and Folder Access Using the Finder
80	Controlling File and Folder Access Using Xsan Admin
81	Unmounting a Volume on a Client
82	Restricting a Client to Read-Only Access
83	Removing a Client from a SAN
84	Mapping Windows User and Group IDs
85	Setting SAN User and Group Quotas
86	About Xsan Quotas
87	Checking User Quota Status
89	Helping Clients Check Their Own Quotas

- 89 Creating Local Home Folders for Network Accounts
- 90 Accessing Client Computers Remotely
 - 90 Controlling a Client Using Screen Sharing
 - 91 Connecting to a Client Using SSH in Terminal
 - 91 Managing a Client Server Using Server Admin

Chapter 6

- 93 **Managing Metadata Controllers**
- 94 Adding a Metadata Controller
- 95 Setting Controller Failover Priority
- 95 Switching to a Standby Controller
- 96 Finding Out Which Controller Is Hosting a Volume
- 97 Listing the Volumes Hosted by a Controller
- 97 Changing a Controller's IP Address
- 98 Accessing Controller Computers Remotely
 - 98 Controlling a Controller Using Screen Sharing
 - 99 Connecting to a Controller Using SSH in Terminal
 - 99 Managing a Controller Using Server Admin
- 99 Monitoring Controller Status

Chapter 7

- 101 **Monitoring SAN Status**
- 101 Checking SAN Status
- 102 Checking Volume Status
- 103 Checking Free Space on a Volume
- 104 Checking Free Space on a Storage Pool
- 104 Graphing SAN Resource Usage
- 105 Setting Up Status Notifications
- 106 Viewing Xsan Logs
- 107 Checking Volume Clients
- 108 Checking for Fibre Channel Connection Failures

Chapter 8

- 109 **Solving SAN Problems**
- 109 If You're Unable to Connect to a Computer Using Xsan Admin
- 109 If You're Unable to Install the Xsan Software
- 109 If Some Computers Aren't Listed in Xsan Admin
- 110 If You're Unable to Mount a Volume on a Client
- 110 If RAID LUNs Aren't Accessible over Fibre Channel
- 110 If Files and Folders Created by Mac OS 9 Computers Show the Wrong Creation Date
- 110 If You Have Problems Using Command-Line Tools
- 110 If a LUN Doesn't Have as Much Space as Expected
- 110 If You're Unable to Rename an Xsan Volume in the Finder
- 111 If You're Unable to Add a Storage Pool
- 111 If Fibre Channel Performance Is Poorer Than Expected
- 112 If a Client is Unable to Use a Volume After a Fibre Channel Interruption

	112	If You're Unable to Add LUNs to a Storage Pool
	113	If the Capacity of a Larger LUN is Listed as 2 Terabytes
	113	If File Copying Doesn't Finish
Appendix A	115	Combining Xsan Controllers and StorNext Clients
	115	Terminology
	116	Compatible Software Versions
	116	Licensing
	117	Using Xsan Controllers with StorNext Clients
Appendix B	119	Using the Command Line
	119	Using the Shell Commands
	119	Sending Commands to Remote Computers
	120	Viewing the Man Pages
	120	Notation Conventions
	120	Installing Xsan from the Command Line
	121	Xsan Commands
	121	Viewing or Changing Volume and Storage Pool Settings (cvadmin)
	124	Manipulating Affinity Tags (cvaffinity)
	124	Copying Files or Folders (cvcp)
	125	Checking or Repairing a Volume (cvfsck)
	126	Labeling, Listing, and Unlabeling LUNs (cvlabel)
	127	Creating a Folder and Assigning an Affinity (cvmkdir)
	127	Creating and Preallocating a File (cvmkfile)
	128	Initializing a Volume (cvmkfs)
	128	Applying Volume Configuration Changes (cvupdatefs)
	128	Defragmenting a File, Folder, or Volume (snfsdefrag)
	130	Controlling the Xsan File System (xsanctl)
	130	Mounting an Xsan Volume
	130	Unmounting an Xsan Volume
	131	Viewing Logs
	131	Xsan Configuration Files
	131	Examples
Glossary	133	
Index	137	

About This Book

Use this guide to learn how to use Xsan 2 set up and manage volumes on a storage area network.

This guide shows how to use Xsan 2 to combine RAID arrays into large, easy-to-expand volumes of storage that clients use like local disks but are actually shared over a high-speed Fibre Channel fabric.

Chapter 1 provides an overview of Xsan and how you can use it to organize RAID arrays into shared volumes of storage.

Chapter 2 describes hardware and software requirements, and offers SAN planning guidelines.

Chapter 3 shows the basic steps for setting up a SAN.

Chapter 4 contains instructions for expanding storage, creating folders with affinities, changing volume and storage pool settings, and checking, defragmenting, and repairing SAN volumes.

Chapter 5 shows how to add client computers to a SAN, mount volumes on clients, control client and user access to SAN files, and control user space through quotas.

Chapter 6 contains information about managing volume metadata controllers.

Chapter 7 shows how to monitor and automatically report the condition of a SAN.

Chapter 8 lists solutions to common problems you might encounter.

Appendix A contains information to help you combine Xsan metadata controllers with Quantum StorNext clients on the same SAN.

Appendix B describes command-line utilities and configuration files you can use to manage an Xsan SAN using the Terminal application.

What's New in Xsan 2

Xsan 2 offers these new features and capabilities:

- The Xsan Admin application has been redesigned to simplify SAN management.
- You can use Xsan Admin to turn on drive activity lights to identify LUNs.
- A volume setup assistant guides you through the process of creating volumes for common purposes such as video editing and file services.
- The volume setup assistant also organizes available storage into storage pools for you, based on the way you plan to use the volume.
- More than one storage pool can have the same affinity tag.
- Each volume has a separate failover priority.

Version Compatibility

The following table shows the compatibility of Xsan 2 metadata controllers and clients with earlier Xsan versions and with StorNext controllers and clients.

Controller	Client	Compatible
Xsan 2	Xsan 2 (Mac OS X v10.5)	Yes
	Xsan 1.4.2 (Mac OS X V10.4 or v10.5)	Yes
	Xsan 1.4–1.4.1	No
	Xsan 1.3 or earlier	No
	StorNext FX 1.4 or 2.0	Yes
	StorNext FS 2.4–3.1	No
Xsan 1.4 or earlier	Xsan 2	No
StorNext FS 3.1	Xsan 2	Yes
StorNext FS 2.4–3.0	Xsan 2	No

Upgrading from an Earlier Version of Xsan

For information about upgrading your SAN storage to Xsan 2 from an earlier version of Xsan, including precautions to take before upgrading and tips for upgrading with the least impact on existing storage, see the *Xsan 2 Migration Guide* at www.apple.com/server/documentation.

Getting Additional Information

For more information about Xsan, consult these resources:

Read Me documents (on the *Xsan Install Disc*)

Read important updates and special information

Server documentation website (www.apple.com/server/documentation)

Get the latest Xsan documentation available in PDF format

Xsan website (www.apple.com/xsan)

Visit the gateway to extensive product and technology information

Xsan Support website (www.apple.com/support/xsan)

Find articles about Xsan from Apple's support organization

Apple Discussions website (discussions.apple.com)

Join a discussion group to share questions, knowledge, and advice with other Xsan administrators

Apple Mailing Lists website (www.lists.apple.com)

Subscribe to mailing lists so you can communicate with other Xsan administrators using email


Notation Conventions

The following conventions are used in this book wherever shell commands or other command-line items are described.

Notation	Indicates
fixed-width font	A command or other text typed in a Terminal window
\$	A shell prompt
[text_in_brackets]	An optional parameter
(one other)	Alternative parameters (type one or the other)
<i>italicized</i>	A parameter you must replace with a value
[...]	A parameter that may be repeated
<angle_brackets>	A displayed value that depends on your SAN configuration

This chapter gives you an overview of Xsan and storage area networks.

Read this chapter for an overview of Xsan and how you can use it to set up a storage area network (SAN) to provide fast, shared storage.



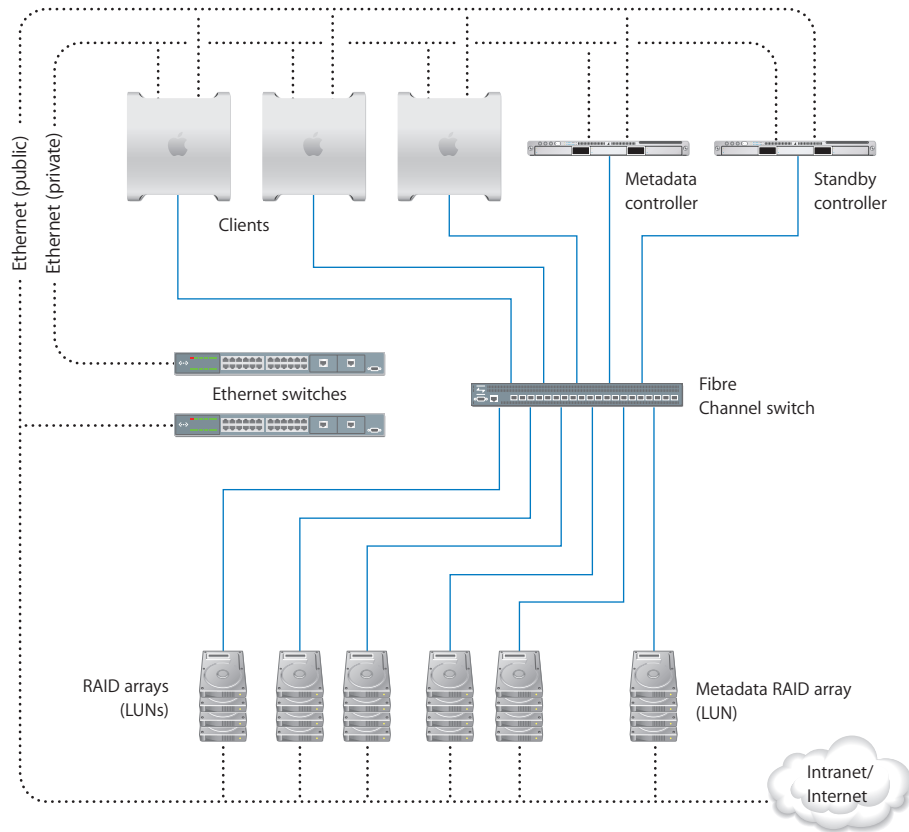
Xsan Storage Area Networks

A storage area network is a way of connecting computers and storage devices so that the computers have fast, shared access to files while making it easy for administrators to expand storage capacity.

An Xsan SAN consists of:

- Shared data volumes
- RAID systems that provide storage space that is protected against disk failure
- At least one computer acting as a metadata controller that combines the RAID arrays and presents their storage to clients as one or more volumes that behave like local disks
- Client computers that access storage in accordance with established permissions and quotas
- Underlying Fibre Channel and Ethernet networks

The following illustration shows the hardware components of an Xsan SAN.



Shared SAN Volumes

A user or application on a client computer accesses shared SAN storage just like they would a local volume. Xsan volumes are logical disks made up of pools of RAID arrays. The elements you combine to create an Xsan volume are described under “How Xsan Storage Is Organized” on page 16.

Metadata Controllers

When you set up an Xsan SAN, you assign at least one computer to act as the metadata controller. The controller manages volume metadata, maintains a file system journal, and controls concurrent access to files. Metadata includes such information as where files are actually stored and what portions of available storage are allocated to new files.

To guarantee volume availability, a SAN should include more than one metadata controller, as shown in the illustration on page 14. If the primary controller fails, the standby controller takes over. Though not recommended for best performance, metadata controllers can also act as clients, so you can use a standby controller as a working client while the primary controller is operational.

Clients

The computers that users or applications use to access SAN volumes are called clients. Clients exchange metadata with controllers over the Ethernet network but use Fibre Channel to send and retrieve file data to and from the RAID systems that provide storage for the volumes.

Network Connections

Xsan uses independent networks to connect storage devices, metadata controllers, and client computers: a Fibre Channel network and one or two Ethernet networks.

Fibre Channel

Xsan moves data between clients and SAN volumes over high-speed Fibre Channel connections. Controllers also use a Fibre Channel connection to move metadata to and from the volume.

Xsan can take advantage of multiple Fibre Channel connections between clients and storage. Xsan can alternate between connections for each read and write, or assign each RAID array in a volume to one of the connections when the volume is mounted.

Ethernet


To prevent metadata traffic from interfering with data on the Fibre Channel network, Xsan controllers and clients exchange file system metadata over a separate Ethernet network. (Controllers do use Fibre Channel to read and write metadata on a volume.) The Xsan Admin application also uses this Ethernet network to let you manage the SAN.

To prevent Internet or intranet traffic from interfering with metadata communications, you should set up separate public (Internet) and private (metadata) Ethernet networks as shown in the illustration on page 14.

How Xsan Storage Is Organized

Although an Xsan volume mounted on a client computer looks like a single disk, it actually consists of multiple physical disks combined on several levels using RAID techniques.

The following illustration shows an example of how disk space provided by the individual drive modules in several RAID systems is combined into a volume that users see as a large local disk.



The following paragraphs describe these elements and how you combine them to create shared Xsan volumes.

LUNs

The smallest storage element you work with in Xsan is a logical storage device called a LUN (a SCSI logical unit number). A LUN represents a group of drives combined into a RAID array.

You create a LUN whenever you create a RAID array on a RAID storage device. The RAID system combines individual drive modules into an array based on the RAID scheme you choose. Each array appears on the Fiber Channel network as a LUN.

Most RAID systems ship already configured as RAID arrays. The corresponding LUNs are ready to use with Xsan.

If the standard RAID arrays on your RAID systems are not right for your application, you can use the RAID system management software to recreate arrays based on other RAID schemes or different numbers of drive modules. For information about other RAID schemes, see “Choosing RAID Schemes for LUNs” on page 30.

The illustration on page 16 shows eight RAID array LUNs. The LUN that stores metadata and journal information uses RAID level 1 (mirrored) to ensure against metadata loss. One LUN stores users’ data on a RAID 0 array (striping only) for best speed and storage efficiency but no data protection. The other data LUNs use RAID 5 (distributed parity) for high performance and storage efficiency with data protection. Xsan sees the RAID arrays as LUNs that can be combined to create a volume.

Your RAID LUNs are labeled and initialized for use with the Xsan file system when you use Xsan Admin to set up a volume.

Storage Pools

LUNs are combined to form storage pools. A storage pool in a small volume might consist of a single RAID array, but a larger volume might consist of several storage pools each of which includes several arrays.

Xsan distributes file data in parallel across the LUNs in a storage pool using a RAID 0 (striping) scheme. So, you can improve a client’s access speed by distributing available storage over several LUNs in a storage pool.

You can set up storage pools that have different performance or recoverability characteristics based on the RAID level of their LUNs, and assign folders to them using affinities. Users can then select where to store files based on their need for speed or safety. See “Folders with Affinities” on page 19.

The illustration on page 16 shows seven LUNs combined into four storage pools for users’ data. One pool uses a single RAID 0 array (fast, but not recoverable). Three other pools use multiple RAID 5 arrays (not as fast, but recoverable), and Xsan stripes data across the LUNs in each of these storage pools.

You use Xsan Admin to add available LUNs to storage pools.

Affinities

Each storage pool is assigned an affinity tag according to the pool’s performance and recoverability characteristics. You can associate a folder with an affinity tag to guarantee that Xsan stores the contents of the folder on a storage pool with the desired characteristics.

More than one storage pool may have the same affinity tag. Xsan distributes the contents of a folder with a particular affinity tag among the storage pools that have that same affinity tag. This strategy improves performance when multiple users simultaneously read and write files in the same folder, because the read and write operations are distributed among the storage pools and their component LUNs.

You use Xsan Admin to assign affinity tags to storage pools and associate folders with those affinity tags.

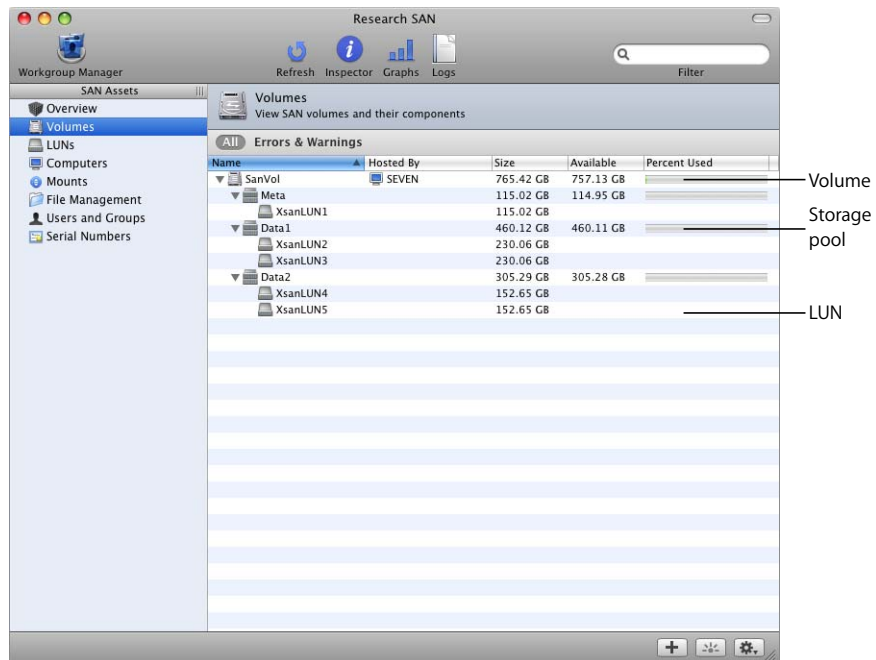
Volumes

Storage pools are combined to create the volumes that users see. From the user's perspective, the SAN volume looks and behaves just like a large local disk, except that:

- The size of the volume can grow as you add underlying arrays or new storage pools
- Multiple users on the SAN can access files on the volume at the same time

In the illustration on page 16, five storage pools are combined to create a single shared volume. You use Xsan Admin to create volumes and mount them on client computers.

The following screen shot shows how LUNs, storage pools, and volumes look as you organize them in Xsan Admin. This example shows a SAN with a single shared volume named "SanVol." Storage for the volume is provided by three storage pools, "Meta," "Data1," and "Data2," the first based on a single LUN and the others on two LUNs each.



Folders with Affinities

To control which storage pools are used to store specific files (for example, to provide different levels of service for different users or applications), you can associate a folder on an Xsan volume with an affinity that is assigned to one or more of the storage pools that make up the volume.

For example, you can associate some folders with an affinity whose storage pools have faster LUNs, and associate other folders with an affinity whose storage pools have safer LUNs. Then users can choose between faster and safer storage by putting files in the appropriate folder.

In the illustration on page 16, the Other folder has an affinity for the faster storage pool that is based on a RAID 0 array. Any file that a user copies into the Other folder is automatically stored on the faster array. The Video and Audio folders are associated with the more secure RAID 5 storage.

How Xsan Uses Available Storage

Xsan stores both user files and file system data on SAN volumes, and stripes data across the LUNs in a volume for better performance.

Metadata and Journal Data

Xsan records information about the files in an Xsan volume using metadata files and file system journals. File system metadata includes information such as which specific parts of which disks are used to store a particular file and whether the file is being accessed. The journal data includes a record of file system transactions that can help ensure the integrity of files in the event of a failure.

These files are managed by the Xsan metadata controller, but are stored on SAN volumes, not on the controller itself. By default, metadata and journal data are stored on the first storage pool you add to a volume. You can use Xsan Admin to choose where these files are stored when you add storage pools to a new volume.

Striping at a Higher Level

When a RAID system writes a file using a RAID 0 (striping) scheme, it breaks the file into segments and spreads them across the individual disk drives in the RAID array. This improves performance by writing parts of the file in parallel (instead of one part at a time) to the individual disks in the array. Xsan applies this same technique at a second, higher level in the storage hierarchy. Within each storage pool in a volume, Xsan stripes file data across the individual LUNs that make up the storage pool. Once again, performance is improved because data is written in parallel.

You can tune SAN performance by adjusting the amount of data written to each LUN in a storage pool (the “stripe breadth”) to suit a critical application.

Security

There are several ways you can control access to a SAN volume:

- Unmount a volume on client computers that shouldn't have access to it. Users cannot browse or mount SAN volumes; only a SAN administrator can mount SAN volumes on clients.
- Mount a volume on a client for read-only access to prevent users on a particular client computer from modifying data on the volume.
- Specify owner, group, and general access permissions in Xsan Admin.
- Specify owner, group, and general access permissions in the Finder.
- Control user access to files and folders on a volume by setting up access control lists (ACLs) in Xsan Admin.
- Set up zones in the underlying Fibre Channel network to segregate users and volumes.

Expanding Storage

There are two ways you can add free space to an existing Xsan volume:

- Add RAID systems (new LUNs) to existing storage pools
- Add entire new storage pools to the volume

Both methods automatically unmount and remount the volume on clients.

You can also add new volumes to a SAN at any time.

For information about expanding Xsan storage, see “Adding Storage” on page 53.

Xsan Capacities

The following table lists limits and capacities for Xsan volumes.

Parameter	Maximum
Number of computers on a SAN (metadata controllers and clients)	64
Number of volumes on a SAN	16
Number of storage pools in a volume	512
Number of LUNs in a storage pool	32
Number of LUNs in a volume	512
Number of files in a volume	4,294,967,296
LUN size	Limited only by the size of the RAID array
Volume size	Limited only by the number and size of LUNs
File size	Approximately 2 ⁶³ bytes
Volume name length	70 characters (A–Z, a–z, 0–9, and _)
File or folder name length	251 ASCII characters
SAN name length	255 Unicode characters
Storage pool name length	255 ASCII characters
Affinity name length	8 ASCII characters
LUN name (label or disk name)	242 ASCII characters

This chapter lists Xsan hardware and software requirements and offers planning guidelines and performance tips that can help you design a SAN that meets your needs.

This chapter contains:

- Xsan hardware and software requirements (page 23)
- SAN planning guidelines (page 27)

Hardware and Software Requirements

Your SAN environment needs to satisfy requirements in these areas:

- Supported computers
- Supported storage devices
- Fibre Channel fabric, adapters, and switches
- Ethernet network
- Directory services (optional)
- Outgoing mail service (optional)

Supported Computers

You can use Xsan 2 on computers that meet these minimum requirements:

Base Systems

- Macintosh computers with an Intel or PowerPC G5 processor

Memory

- Client computers need at least 2 GB of RAM (clients with Xsan 1.4.2 and Mac OS X v10.4 Tiger or Mac OS X Server v10.4 Tiger need at least 1 GB).
- Computers used as metadata controllers need at least 2 GB of RAM for a single volume plus an additional 2 GB of RAM for each additional SAN volume hosted by the controller.

Supported Operating Systems

You can install Xsan 2 only on computers with Mac OS X v10.5 Leopard or Mac OS X Server v10.5 Leopard.

Mac client computers that have Xsan 1.4.2 can join an Xsan 2 SAN. These clients can have Leopard or Tiger.

To join an Xsan 2 SAN, Windows, AIX, IRIX, Linux, and Solaris clients must be running Quantum's StorNext File System version 2.6 or 2.7. For complete compatibility information, see "Version Compatibility" on page 10.

Supported Storage Devices

Although you can use any standard SCSI LUN storage device, this guide assumes you are using Apple-approved RAID systems for your storage devices. For the latest information about qualified RAID systems, see the Xsan webpage at:

www.apple.com/xsan

Important: Be sure to install the latest firmware update on your RAID systems before you use them with Xsan.

Fibre Channel Fabric

Unlike file system metadata, which controllers and clients exchange over Ethernet, actual file content in an Xsan SAN is transferred over Fibre Channel connections (as is metadata that controllers access on a volume). To set up the connections, you need:

- An Apple Fibre Channel PCI, PCI-X, or PCI-E card for each client and controller computer
- A supported Fibre Channel switch
- Fibre Channel cables connecting computers and storage devices to the switches to form a Fibre Channel fabric

Fibre Channel PCI Cards

Install Apple Fibre Channel PCI, PCI-X, or PCI-E cards in all Macintosh computers that will connect to the SAN.

Fibre Channel Switches

Fibre Channel switches from Brocade, Cisco, and QLogic have been tested with Xsan and the Apple Fibre Channel PCI, PCI-X, and PCI-E cards. For the latest information about qualified switches, see the Xsan webpage at:

www.apple.com/xsan

Fabric Configuration

You must connect the computers, storage devices, and switches in your Fibre Channel network to form a Fibre Channel “fabric.” All Apple-approved switches create a fabric by default as soon as you plug in the Fibre Channel cables. In a fabric, Fibre Channel cables connect node ports (F or N_Port). See the documentation that came with your Fibre Channel switches for more information.

Ethernet TCP/IP Network

Computers on the SAN must also be connected to an Ethernet network. Xsan controllers and clients use this network instead of the Fibre Channel network to exchange file system metadata.

If the computers on your SAN need to communicate with directory servers, a corporate or campus intranet, or the Internet, you should connect each SAN client and metadata controller to two separate Ethernet networks: one private subnet for the SAN metadata and a separate connection for directory service, intranet, and Internet traffic. This is especially important if you plan to use the SAN for high-performance applications such as video editing.

IP Addresses

The client and metadata controller computers need static (fixed) IP addresses for their Ethernet network connections. For the public intranet and Internet connection, you can enter each computer’s static IP address, subnet mask, router address, and DNS server address manually or configure a DHCP server to provide some or all of this information. If you want the DHCP server to provide IP addresses, it must always assign the same static IP address to each SAN computer. Don’t use DHCP to assign dynamic IP addresses to SAN devices.

For the SAN metadata network, the SAN computers should have static private (non-routable) IP addresses (unless you’re unable to set up a separate, private Ethernet network for SAN metadata). If you’re setting up new computers or computers on which you have just performed a clean installation of Leopard or Leopard Server, you can have Xsan Admin assign and manage addresses for your private metadata network. If you choose to assign addresses yourself, you can use one of the following ranges of IP addresses on your private (non-routed) metadata network:

Private address range	Associated subnet mask	Comments
10.0.0.0 – 10.255.255.255	255.0.0.0	10/8
172.16.0.0 – 172.31.255.255	255.240.0.0	172.16/12
192.168.0.0 – 192.168.255.255	255.255.0.0	192.168/16

Directory Services

If you plan to use user and group privileges to control access to files and folders on the SAN, you should set up or join a central directory of users and groups. A central directory service lets you manage all SAN users and groups from one computer instead of having to visit and painstakingly configure each SAN client and metadata controller.

If you already have directory service provided by an Open Directory server, you can have the setup assistant configure each metadata controller and client computer with Xsan 2 to use existing user and group accounts from the Open Directory server. If you have another type of directory service, such as Active Directory, you configure each controller and client to connect to it for user and group accounts by using the Directory Utility application after initial setup. If you have client computers with Tiger, you use the Directory Access application on each one to connect it to a directory server.

If your SAN doesn't have access to an existing directory service, you can specify during initial setup of your Xsan primary metadata controller that you want to use Xsan Admin to manage your users and groups. The setup assistant creates an Open Directory master server on your primary metadata controller and sets up Open Directory replica servers on your standby metadata controllers. The Open Directory master provides an LDAP directory, single sign-on user authentication using Kerberos, and password validation using common authentication methods. The replicas improve responsiveness and provide automatic failover of Open Directory services.

The setup assistant also configures client computers that have Xsan 2 installed to connect to your Xsan primary metadata controller for Open Directory user and group accounts. If you have client computers with Tiger, you need to use the Directory Access application on each one to connect it to your Xsan primary metadata controller's Open Directory service.

If you need to set up an Open Directory server yourself, you can use Mac OS X Server's Server Admin application. Then you use the Workgroup Manager application to manage users and groups. For information, see *Open Directory Administration* and *User Management* at www.apple.com/server/documentation.

Note: Some applications running on SAN client computers, such as Final Cut Pro, work better when users have local home folders, not network home folders. User accounts that you manage with Xsan Admin are automatically set up with local home folders. For help setting up local home folders for user accounts that you don't manage with Xsan Admin, see "Creating Local Home Folders for Network Accounts" on page 89.

If you decide not to use a central directory service, you need to set up the same users and groups in the Accounts pane of System Preferences on each SAN computer.

Important: If you create users and groups on each SAN computer, be sure that:

- Each user or group has a numeric user ID (UID) or group ID (GID) that is unique throughout the SAN
- Each user or group defined on more than one computer has the same UID or GID on each computer

Outgoing Mail Service

Xsan can send SAN status notifications via email on your local network (IP subnet) without using a separate mail server. However, to send notifications outside your local network, you need an SMTP server to act as a mail gateway. If you don't have access to an outgoing mail server, you can use the mail service in Mac OS X Server to set one up. For information, see *Mail Service Administration* at www.apple.com/server/documentation.

Planning Your SAN

It's easy to add storage to an existing Xsan SAN, but reorganizing a SAN after you set it up is not so simple. So, it's important to plan the layout and organization of your SAN and its storage before you set it up.

An Xsan SAN is made up of:

- Storage devices (RAID systems)
- LUNs (SCSI logical unit numbers, usually RAID arrays)
- Storage pools (groups of LUNs)
- Affinity tags, which identify storage pools with similar performance and data protection
- Volumes (groups of storage pools visible to users)
- Clients (computers that use volumes)
- Controllers (computers that manage volume metadata)
- An Ethernet network used to exchange volume metadata
- A Fibre Channel network used to transfer data to and from volumes

Before you try to set up a SAN, you need to decide how you want to organize these components. Take the time to create a diagram or a table that organizes available hardware into RAID arrays, volumes, client computers, and metadata controllers in a way that meets SAN users' needs and your needs as the SAN administrator. You don't need to plan your storage pools or affinity tags if you set up each volume using a preset volume type based on the kind of work the volume will support.

Preliminary Planning Questions

As you plan, consider the following questions:

- How much storage do you need?

- How do you want to present available storage to users?
- What storage organization makes the most sense for user workflow?
- What levels of performance do your users require?
- How important is high availability?
- What are your requirements for security?

Your answers to the above questions will help you decide the following:

- What RAID schemes should you use for your RAID arrays?
- How many SAN volumes do you need?
- How should individual volumes be organized?
- Which preset volume type can you choose for each volume?
- Which LUNs should be assigned to each affinity tag?
- Which clients, users, and groups should have access to each volume?
- Which computers will act as metadata controllers?
- Do you need standby metadata controllers?
- Do you want to use metadata controllers as clients also?
- Do you need to customize the storage location of file system metadata and journal data?
- Do you need to adjust a volume's allocation strategy?
- How should you configure your Ethernet network?

Review the considerations and guidelines on the following pages for help translating your answers into a suitable SAN design.

Planning Considerations and Guidelines

The following paragraphs might help you make some of your SAN design decisions.

How Much Storage?

Because it's easy to add storage for user data to an Xsan SAN, you only need to decide on an adequate starting point. You can add storage later as needed.

However, you can't expand a storage pool that can only store volume metadata and journal data, so you should try to allocate enough space for metadata right from the start. (You *can* add an entire storage pool for metadata and journal storage.) For help estimating your metadata and journal data storage requirements, see "Estimating Metadata and Journal Data Storage Needs" on page 34.

Note that the number of RAID systems you use affects not only available space but also SAN performance. See "Performance Considerations," below.

How Should Users See Available Storage?

If you want the users working on a particular project to see a volume dedicated to their work, create a separate volume for each project. If it's acceptable for a user to see a folder for his or her work on a volume with other peoples' folders, you can create a single volume and organize it into project folders.

Workflow Considerations

How much file sharing is required by your users' workflow? If, for example, different users or groups work on the same files, either simultaneously or in sequence, it makes sense to store those files on a single volume to avoid having to maintain or hand off copies. Xsan uses file locking to manage shared access to a single copy of the files.

Performance Considerations

If your SAN supports an application (such as high resolution video capture and playback) that requires the fastest possible sustained data transfers, design your SAN with these performance considerations in mind:

- Set up the LUNs (RAID arrays) using a RAID scheme that offers high performance. See "Choosing RAID Schemes for LUNs" on page 30.
- Assign your fastest LUNs to an affinity tag for the application. Assign slower LUNs to an affinity tag for less demanding applications.
- To increase parallelism, spread LUNs across different RAID controllers. Xsan then stripes data across the LUNs and benefits from simultaneous transfers through two RAID controllers.
- To increase parallelism for an affinity tag assigned to relatively small LUNs (the size of one or a few drive modules), create a slice of similar size across all the drives on a RAID controller instead of creating the LUNs from just one or two drive modules.
- Spread file transfers across as many drives and RAID controllers as possible. Try creating slices across the drives in RAID systems, and then assign these slices to the same affinity tag.
- To increase throughput, connect both ports on client Fibre Channel cards to the fabric.
- Store file system metadata and journal data on a separate storage pool from user data, and make sure the metadata LUNs are not on the same RAID controller as any user data LUNs.
- Use a second Ethernet network (including a second Ethernet port in each SAN computer) for the SAN metadata, or at least use a router to isolate the Ethernet network used by the SAN from a company intranet or the Internet.
- If your SAN uses directory services, mail services, or other services on a separate server, connect SAN computers to that server on an Ethernet network separate from the SAN metadata network.

- Choose a different primary metadata controller for each volume, and set up volume failover priorities to minimize the possibility of more than one volume failing over to the same metadata controller.

Availability Considerations

If high availability is important for your data, set up at least one standby metadata controller in addition to your primary metadata controller. Also, consider setting up dual Fibre Channel connections between each client, metadata controller, and storage device using redundant Fibre Channel switches.

WARNING: Losing a metadata controller without a standby can result in the loss of all data on a volume. A standby controller is recommended.

Security Considerations

If your SAN will support projects that need to be completely secure and isolated from each other, you can create separate volumes for each project to eliminate any possibility of the wrong client or user accessing files stored on a volume.

As the SAN administrator, you control which client computers can use a volume. Clients can't browse for or mount SAN volumes on their own. You use Xsan Admin to unmount a volume on clients that shouldn't have access to it.

You can also set up access control lists (ACLs) in Xsan Admin or assign user and group permissions to folders using standard file access permissions in the Finder.

Choosing RAID Schemes for LUNs

Much of the reliability and recoverability of data on a SAN is not provided by Xsan itself but by the RAID arrays you combine to create your storage pools and volumes. Before you set up a SAN, you use the RAID system configuration or administration application to prepare LUNs based on specific RAID schemes.

WARNING: If a LUN belonging to an Xsan volume fails and can't be recovered, all data on the volume is lost. It is strongly recommended that you use only redundant LUNs (LUNs based on RAID schemes other than RAID 0) to create your Xsan volumes.

LUNs configured as RAID 0 arrays (striping only) or LUNs based on single drives are difficult or impossible to recover if they fail. Unprotected LUNs such as these should only be used for volumes that contain scratch files or other data that you can afford to lose.

Most RAID systems support all popular RAID levels. Each RAID scheme offers a different balance of performance, data protection, and storage efficiency, as summarized in the following table.

RAID level	Storage efficiency	Read performance	Write performance	Data protection
RAID 0	Highest	Very High	Highest	No
RAID 1	Low	High	Medium	Yes
RAID 3	High to very high	Medium	Medium	Yes
RAID 5	High to very high	High	High	Yes
RAID 0+1	Low	High	High	Yes

Deciding on the Number of Volumes

A volume is the largest unit of shared storage on the SAN. If your users need shared access to files, you should store those files on the same volume. This makes it unnecessary for them to pass copies of the files among themselves.

On the other hand, if security is critical, one way to control client access is to create separate volumes and unmount volumes on clients that shouldn't have access to them.

For a more typical balance of security and shared access, a flexible compromise is to create a single volume and control access with folder access privileges or access control lists in Xsan Admin (or in Mac OS X Server's Server Admin).

Deciding How to Organize a Volume

You can help users organize data on a volume or restrict users to specific areas of the volume by creating predefined folders. You can control access to these folders by assigning access permissions using Xsan Admin.

You can assign folders to specific storage pools using affinities. You can, for example, create a folder for data that requires fast access and assign that folder to your fastest storage pool.

Assigning LUNs to Affinity Tags

When you create a volume using a preset volume type that fits your SAN scenario, Xsan Admin automatically sets up storage pools and affinity tags for best performance. All you do is assign LUNs to each affinity tag. Xsan Admin determines the optimal number of storage pools to create, based on the volume type and the number of LUNs you assign to each affinity tag. For best performance, you should assign LUNs in the multiples shown below. These multiples apply to affinity tags used for user data, not to the Metadata and Journal affinity tag, which needs just one LUN.

For this volume type's affinity tags used for user data	Assign LUNs in multiples of
Calendar Server Cluster	1
General File Server	2
Home Folder Server	2
Mail Cluster	1
Podcast Producer Cluster	4
Standard Definition Video	4
Uncompressed High Definition Video	4

You should assign LUNs that have the same capacity and performance characteristics to each affinity tag.

LUNs that you assign to an affinity tag should have the same capacity because Xsan provides high performance by using the RAID 0 scheme to stripe data across the LUNs in each storage pool. This striping scheme can use available space on each LUN equal to the capacity of the smallest LUN in a storage pool. If a storage pool's LUNs vary in size, this can result in wasted capacity. For example, if a storage pool has a 240 GB RAID array and a 360 GB RAID array, 120 GB of the larger array will not be used. By assigning LUNs with similar capacities to an affinity tag, you avoid wasting available storage.

If you're using a volume type with multiple affinity tags for user data, assign your fastest LUNs to the affinity tag that will be associated with folders whose contents will benefit most from extra performance. Assign slower LUNs to an affinity tag that will be associated with folders whose contents don't have critical performance requirements.

You can also increase the performance of an affinity tag's storage pools by assigning that affinity tag a combination of LUNs that are hosted on different drive modules and different RAID controllers. This strategy increases performance by increasing the parallelism of data transfers.

Deciding Which Clients to Mount a Volume On

If you create multiple volumes, decide which volumes should be mounted on which clients. A new volume is mounted initially on all clients. You can use Xsan Admin to unmount a volume from selected clients.

Choosing Metadata Controllers

You must choose at least one computer to be the SAN metadata controller, the computer that is responsible for managing file system metadata.

Note: File system metadata and journal data are stored on the SAN volume, not on the metadata controller itself. For more information, see “Choosing Where to Store Metadata and Journal Data” on page 33.

If you have a small number of clients or if performance is not critical, you can use a single computer as both client and metadata controller. You can even set up a SAN consisting of a single storage device and a single computer that acts as both controller and client (to provide network attached storage, for example).

If high availability is important, you should use at least two metadata controllers, one as the primary controller and one as a standby. You can specify additional metadata controllers as needed, and set each volume’s failover priorities to determine the order in which the controllers are tried if a volume’s primary controller stops responding.

If performance is critical, don’t run other server services on the metadata controller and don’t use the controller itself to reshare a SAN volume using AFP or NFS.

Choosing Standby Controllers

To be sure that SAN volumes are always available, set up at least one standby metadata controller that can take over if your primary metadata controller fails.

Combining Clients and Controllers

The same computer can function as both a metadata controller and a client. It’s possible, for example, to set up a SAN consisting of a single RAID system and one computer that acts as both controller and client. Any computer you specify as a controller can also act as a client.

If, for example, you don’t have a computer to dedicate as a standby metadata controller, you can assign a computer that is normally used as a client to take over controller duties if the primary metadata controller fails.

To keep clients and controllers separate, you can set up client-only computers for your users.

Choosing Where to Store Metadata and Journal Data

The metadata and journal data that describe a volume are not stored on the volume’s metadata controller but on the volume itself. By default, they are stored on the first storage pool in the volume.

All preset volume types set up a separate storage pool used only for metadata and journal data. For best performance, make sure that the LUNs you assign to the metadata and journal storage pool are connected to a different RAID controller than the LUNs that you assign to affinity tags for user data.

If you set up a custom volume with more than one storage pool, you can choose which storage pool is used to store metadata and journal data. You may get adequate performance by combining metadata and journal data on the same storage pool as user data, but for better performance, use a separate storage pool for metadata and journal data.

Estimating Metadata and Journal Data Storage Needs

To estimate the amount of space required for Xsan volume metadata, assume that 10 million files on a volume require approximately 10 gigabytes of metadata on the volume's metadata storage pool.

Choosing an Allocation Strategy

If you choose a preset volume type when you set up a volume, Xsan Admin sets its volume allocation strategy for you. Later, you can change the allocation strategy by editing volume settings with Xsan Admin. The allocation strategy you choose for a volume determines the order in which its storage pools are filled with data. You can choose *round robin*, *fill*, or *balance*.

If you choose round robin, Xsan writes new data to each storage pool in the volume in turn. This is normally the best choice for performance.

If you choose fill, Xsan writes all new data to the first storage pool in the volume until that storage pool is full, and then moves to the next storage pool. This is a good choice if you want to keep a particular storage pool unused as long as possible.

If you choose balance, Xsan writes new data to the storage pool with the most free space.

Planning the Ethernet TCP/IP Network

Ethernet connections are used in several ways in an Xsan storage area network:

- Xsan clients and metadata controllers use Ethernet to exchange volume metadata.
- Xsan clients can use Ethernet for access to networks outside the SAN (campus or corporate intranet or the Internet).
- Xsan metadata controllers can use Ethernet connections for remote management.
- RAID systems can use Ethernet connections for system management.
- Fibre Channel switches can use Ethernet connections for switch management.

You have two options:

- Use one Ethernet network for all traffic. This is the less expensive option, but is also less secure and might not provide the best possible performance.

- Use two separate networks—one for metadata and another for all other IP traffic. This configuration is slightly more expensive (requiring two Ethernet adapters for each computer) but offers greater security and better performance because routine network traffic doesn't interfere with SAN volume metadata traffic.

Using a Private Metadata Network

Non-SAN-related Ethernet traffic can interfere with the exchange of metadata among Xsan controllers and clients. For example, using the same connection for both Xsan metadata exchange and Internet access can slow file system performance. Similarly, using the same Ethernet network to connect client computers to directory services and SAN metadata can affect SAN performance.

If SAN performance is critical for your users or applications, keep all extraneous traffic off the network that clients and metadata controllers use to exchange metadata. For best SAN performance, set up a private Ethernet TCP/IP network for the exclusive use of Xsan clients and metadata controllers. For other types of network traffic, including Internet access, RAID system and Fibre Channel switch management, remote SAN management, or directory services, connect each client and metadata controller to a second Ethernet network using a second Ethernet port.

Using Switches Instead of Hubs

To get the best performance, use Ethernet switches, not hubs, in the SAN metadata network.

Planning the Fibre Channel Network

Xsan uses Fibre Channel connections to:

- Transfer user data directly between clients and data storage pools
- Transfer metadata between metadata controllers and metadata storage pools

You should verify Fibre Channel performance and troubleshoot the fabric if you have connections operating below 2 GB/s.

Verifying Base Fibre Channel Performance

Because the devices connected to a Fibre Channel network automatically adjust their speed to match the slowest device on the fabric, it is important to check that all connections in the fabric are operating at 2 GB/s.

To check Fibre Channel connection performance:

- Use the management software provided with your Fibre Channel switches to test the performance of your Fibre Channel fabric.

If Your Fibre Channel Fabric Is Running Slower Than Expected

The following paragraphs list things you can check if your Fibre Channel fabric is not running at the expected 2 GB/s.

Check Cables

One faulty cable in a fabric can slow the entire network. Check all cables to make sure they are capable of full transmission speed. Use your switch management software to isolate the faulty cable by checking the performance of specific connections.

Use Qualified Transceivers in Matching Pairs

Check with the manufacturers of the devices you are connecting to your fabric to be sure that the transceivers (GBICs) you are using are qualified for use with their devices.

Also, use identical transceivers (same manufacturer and model number) on both ends of each cable. Mismatched optical transceivers (even if they are both separately qualified for use with your devices) can cause Fibre Channel communication errors and degrade SAN performance.

Check Fibre Channel Switch Port Configuration

The Request for State Change Notifications (RSCN) that is generated when a client on the SAN restarts can cause dropped frames in video streams to other clients.

To avoid interrupting SAN traffic to other clients if one client restarts, check your Fibre Channel switch documentation to see if you can configure the switch to suppress RSCNs on initiator ports. (On Qlogic switches, for example, this feature is called I/O StreamGuard.)

Connect Devices to Specific Blades

If your Fibre Channel switch is based on a blade architecture, you might be able to improve performance by:

- Connecting pairs of devices that routinely exchange large volumes of data to the same blade in the switch
- Distributing loads across multiple blades instead of concentrating all of the load on one or two blades

Configuring RAID Systems

Follow these guidelines when you set up your RAID systems for use as Xsan LUNs.

Installing the Latest Firmware

To get the best performance and reliability from your RAID systems, be sure to install the latest available firmware.

Connecting RAID Systems to an Ethernet Network

For best performance, don't connect RAID controller Ethernet management ports to the SAN's metadata network. Connect the ports to the separate Ethernet network that you use for other types of network traffic, such as directory services, Internet access, and remote Xsan management.

Choosing RAID Levels for LUNs

Use RAID 1 for metadata LUNs and RAID 5 for data LUNs.

Use RAID 1 for Metadata LUNs

RAID 1 (mirroring) can give slightly better performance than the default RAID 5 scheme for the small, two-drive metadata LUNs that Xsan uses to store volume information. A single drive is almost always adequate for storing the primary volume metadata (10 GB of metadata space is enough for approximately 10 million files). The second, mirror drive protects you against metadata loss.

Use RAID 5 for Data LUNs

Most RAID systems are optimized for excellent performance and data redundancy using a RAID 5 scheme. (RAID 5 stripes data across the available drives and also distributes parity data across the drives.) Some RAID systems ship already configured as RAID 5 LUNs. RAID 0 (striping with no parity) might give slightly better write performance but provides no data recovery protection, so RAID 5 is always a better choice for LUNs used to store user data.

Adjusting RAID System Performance Settings

RAID system performance settings, which affect parameters such as drive caching, RAID controller caching, and read prefetching, can have a significant effect on Xsan volume performance. Follow these guidelines.

Enable Drive Caching

In addition to the caching performed by the RAID controller, each drive in an array can perform its own caching at the drive level to improve performance.

WARNING: If you enable drive caching for a RAID set, be sure that the system is connected to a UPS. Otherwise, you could lose cached data if the power fails.

Enable RAID Controller Write Caching

Without RAID controller write caching, a request to write data to the associated LUN is not considered finished until the data has been completely written to the physical disks that make up the array. Only then can the next write request be processed. (This is sometimes called "write-through caching.")

When RAID controller write caching is enabled, a request to write data is considered finished as soon as the data is in the cache. This is sometimes called “write-back caching.” Write requests are processed more quickly because the file system just needs to write to the fast cache memory and doesn’t need to wait for the slower disk drives.

Be sure to enable write caching on RAID controllers that support metadata storage pools.

Although some large write requests might benefit from caching, often they don’t. By placing a volume’s metadata storage pool on a RAID controller separate from the data storage pools, you can enable caching on the RAID controller used for metadata, and disable caching on the RAID controller used for data.

When the file system is relying on caching in this way, you must guarantee that data in the cache isn’t lost before it is actually written to disk. Data that has been written to disk is safe if the power fails, but data in a cache is not. To be sure that a power failure can’t cause the loss of cached data, protect your RAID systems with RAID controller backup batteries or an uninterruptible power supply (UPS).

WARNING: If you enable controller write caching on a RAID system, be sure that the system includes controller backup batteries and, preferably, is connected to a UPS.

Enable Read Prefetching

Read prefetching is a technique that improves file system read performance in cases where data is being read sequentially, as in the case of audio or video streaming, for example. When read prefetching is enabled, the RAID controller assumes that a read request for a particular block of data will be followed by requests for subsequent, adjacent data blocks. To prepare for these requests, the RAID controller reads not only the requested data, but also the following data, and stores it in cache memory. Then, if the data is actually requested, it is retrieved from the fast cache instead of from the slow disk drives.

Setting Up a Storage Area Network

3

This chapter shows step by step how to set up a shared volume on an Xsan SAN.

This chapter explains how to connect the SAN networks, prepare RAID arrays (LUNs), use the Xsan Admin application, set up a new SAN, and create a shared volume. This chapter also tells you how to administer Xsan remotely, rename a SAN, remove a SAN, set up additional SANs, and manage multiple SANs.

Connecting Computers and Storage Devices

Before you open Xsan Admin to configure your SAN, you need to connect client computers, controller computers, and storage devices to the SAN's Fibre Channel and Ethernet networks. Make sure your networks meet the requirements summarized under "Fibre Channel Fabric" on page 24 and "Ethernet TCP/IP Network" on page 25.

Preparing LUNs

New RAID systems often come already configured as one or more RAID arrays. So, out of the box, your RAID system might provide LUNs that you can use for most SAN applications. For details, see the documentation for your RAID system. Unless you have well-defined, special needs, no other LUN preparation is needed.

If you want to set up some other combination of RAID arrays or slices, use the management application that comes with your RAID systems to create other arrays before you add the resulting LUNs to your SAN's storage pools. For information about choosing a RAID scheme, see "Choosing RAID Schemes for LUNs" on page 30.

Note: Don't use Disk Utility to format arrays or slices for use with Xsan. LUNs are labeled and initialized when you add them to a storage pool using Xsan Admin. After they are labeled, the LUNs can't be modified using Disk Utility.

Be sure to create arrays of the same size if you plan to add them to the same storage pool. For more information, see "Assigning LUNs to Affinity Tags" on page 32.

Using Server Assistant to Configure Controllers

You can use the Mac OS X Server setup assistant, Server Assistant, to configure servers as Xsan metadata controllers. Server Assistant runs automatically when you start up a new server or a server on which you have performed a clean installation of Mac OS X Server v10.5 Leopard. When using Server Assistant to set up your primary metadata controller, you can choose how you will manage Xsan users and groups.

Managing Users and Groups with Xsan Admin

When using Server Assistant to set up your primary metadata controller, you can choose to manage SAN users and groups with Xsan Admin. This option is recommended if you don't already have a directory server and you expect to have up to 20 SAN users. These users will have local home folders on their computers (not network home folders on the server).

Important: You can choose the option to manage users and groups with Xsan Admin only when you use Server Assistant to set up your primary metadata controller. You can't configure this option after setting up the primary controller with Server Assistant.

If you choose this option, Server Assistant makes the primary metadata controller an Open Directory master. Then Xsan Admin automatically configures standby metadata controllers as Open Directory replica servers. For the Open Directory master and replicas, the directory administrator's user name is Directory Administrator, the short name is diradmin, and the password is initially the same as the password of the administrator account that you create with Server Assistant.

Xsan Admin also automatically configures client computers with Xsan 2 to connect to your Xsan primary metadata controller for Open Directory user and group accounts. If you have client computers with Tiger, you need to use the Directory Access application on each one to connect it to your Xsan primary controller's Open Directory service.

Using an Existing Open Directory Server

If you have an existing Open Directory server, you can choose to have Server Assistant configure your primary metadata controller to get users and groups from it. Then when you set up your SAN, the Xsan setup assistant automatically configures standby metadata controllers and client computers with Xsan 2 to connect to the Open Directory server for users and groups.

Using Another Directory Server

When using Server Assistant to set up your primary metadata controller, you can also choose to connect to an existing directory server later, including Active Directory or Open Directory, after you finish setup. In this case, you use Directory Utility (in /Applications/Utilities/) on each metadata controller and client computer to configure a connection to your directory server. If you have client computers with Tiger, use Directory Access to connect them to your directory server.)

Using Xsan Admin

You use the Xsan Admin application (installed in /Applications/Server/) to set up and manage your SAN. You can use Xsan Admin to manage an Xsan 2 SAN from any computer that has access to the same public intranet as the SAN.

Installing Just the Xsan Admin Application

You can install Xsan Admin on any computer you want to use to manage the SAN. See “Setting Up an Xsan Administrator Computer” on page 51.

Connecting Through a Firewall

If there is a firewall between the SAN and the computer you’re running Xsan Admin on, be sure port 311 in the firewall is open so Xsan Admin can communicate with the SAN computers.

Xsan Admin Preferences

Open Xsan Admin and choose Xsan Admin > Preferences to adjust these settings:

- SAN status refresh interval
- The amount of log information displayed
- The maximum number of users to list when searching

Getting Help

Xsan Admin includes onscreen help. Use the Help menu, or click the help button in any Xsan Admin dialog or pane.

SAN and Volume Setup Summary

To set up a shared volume on a SAN, you’ll perform the following tasks.

- 1 Set up the Fibre Channel network (page 42)
- 2 Set up the Ethernet network (page 42)
- 3 Configure computers to use a network time server (page 43)
- 4 Set up SAN users and groups (page 43)
- 5 Set up RAID systems (page 44)
- 6 Set up a metadata array (page 45)
- 7 Install Xsan software on SAN computers (page 45)
- 8 Configure the SAN (page 47)
- 9 Create a volume (page 49)
- 10 (Optional) Set up status notifications (page 51)
- 11 (Optional) Assign folders to affinity tags (page 51)
- 12 (Optional) Set user and group quotas (page 51)

Setting Up an Xsan Volume on a Storage Area Network

Step 1: Set Up the Fibre Channel Network

- Connect controller computers, client computers, and RAID storage systems to a Fibre Channel switch to create a Fibre Channel fabric for the SAN. Be sure to configure the switch and make the connections so that you create a Fibre Channel fabric.

For more information, see the guidelines and requirements under “Fibre Channel Fabric” on page 24.

Step 2: Set Up the Ethernet Networks

- 1 Connect controller computers, client computers, and RAID systems to the public intranet and Internet.
- 2 If you’re going to use a private metadata network for the performance benefit it provides, set up its Ethernet switch and connect the SAN computers to it.
- 3 Configure the network settings on the client and controller computers.

For each computer’s public Ethernet port, you can configure the TCP/IP settings:

- **Manually:** You enter the static IP address, subnet mask, router address, and DNS server address for each computer.
- **Using DHCP with manual address:** You enter the computer’s static IP address, and your DHCP server provides the other TCP/IP connection settings.
- **Using DHCP:** Your DHCP server provides a static IP address and the other TCP/IP settings for client computers. (This configuration method isn’t available when setting up metadata controllers.) The DHCP server must be configured to always assign the same static IP address to each SAN computer.

For the private metadata network, you can have the Xsan setup assistant configure the network settings if you’re setting up new computers or computers on which you have just performed a clean installation of Leopard or Leopard Server. To make sure the Xsan setup assistant offers this option, don’t configure the Ethernet port connected to the private metadata network:

- On client computers with Leopard, leave this Ethernet port unconfigured in Network preferences.
- On metadata controllers and clients with Leopard Server, disable this Ethernet port while using the Mac OS X Server setup assistant.

The Xsan setup assistant offers to configure the private metadata network if it finds exactly one available unconfigured Ethernet port on each computer, or if each computer already has an Ethernet port with a private IP address on the same IP subnet and a subnet mask of 255.255.255.0.

For information about private IP addresses and the network settings you need to make if the Xsan setup assistant doesn’t configure the metadata network settings on SAN computers, see “Ethernet TCP/IP Network” on page 25.

Step 3: Configure SAN Computers to Use a Time Server

- To ensure consistent time metadata across all computers in the SAN, open Date & Time preferences and choose the same network time server for all metadata controller and client computers.

If you're setting up a new server or a computer on which you have performed a clean installation of Leopard Server, you can choose a network time server in the Time Zone pane of the Mac OS X Server setup assistant.

Step 4: Set Up SAN Users and Groups

Here are several ways you can set up users and groups for your SAN:

- If you're setting up a new primary metadata controller or one on which you have just performed a clean installation of Leopard Server, select an option in the "Users and Groups" pane of the Mac OS X Server setup assistant.



- **Manage users and groups with Xsan Admin:** Select this option to have the server setup assistant create a centralized directory of users and groups on the primary metadata controller.

Important: You can select this option only while setting up Leopard Server on the primary metadata controller. You can't configure this option after using the Mac OS X Server setup assistant on the primary controller.

Later, in the “Add User Accounts” pane, you have the option of creating some user accounts. The Mac OS X Server setup assistant creates the user accounts in the LDAP directory of the Open Directory server it creates on your primary metadata controller. After setup, you use Xsan Admin to create and delete users and groups and to change group membership. For information about Open Directory servers, see “Directory Services” on page 26.

- **Use existing users and groups from an Open Directory server:** Select this option to have the Mac OS X Server setup assistant configure the primary metadata controller to connect to the Open Directory server whose DNS name or IP address you specify. If you select this option, Xsan Admin automatically configures all other SAN computers with Xsan 2 to use the Open Directory server.

If you need to set up an Open Directory domain, you can use Mac OS X Server’s Server Admin application. Then you use the Workgroup Manager application to manage users and groups. For information, see *Open Directory Administration* and *User Management* at www.apple.com/server/documentation.

- **Connect to a directory server later:** Select this option if you have another type of directory server, such as Active Directory.
After setting up the primary metadata controller, use the Directory Utility application on each SAN computer with Leopard or Leopard Server to connect the computer to your directory server. If you have client computers with Mac OS X v10.4 Tiger or Mac OS X Server v10.4 Tiger, use the Directory Access application to connect these computers to the directory of users and groups.

- If you choose not to use a directory service, you need to create the same set of users and groups in System Preferences on each SAN computer.

Important: If you create users and groups individually on each SAN computer, be sure that each user and group has a numeric user ID (UID) or group ID (GID) that is unique throughout the SAN, and all SAN users and groups have the same UID or GID on all SAN computers. One way to do this is to create an identical list of users and groups in the same order on each computer following a clean installation of the operating system.

Step 5: Set Up RAID Systems

- 1 Follow the instructions that come with your RAID systems to turn them on and configure their network, management, and security settings.
- 2 If your RAID systems come with RAID sets already configured, they’ll be automatically detected during SAN setup, and you can skip to Step 6, “Create a Metadata Array.” Otherwise, use the management software that comes with the RAID system to create arrays based on the RAID schemes of your choice, leaving three drives on one system unassigned so you can create a small, separate metadata LUN as described in the next step.

New RAID systems are usually preconfigured as a set of RAID arrays that are ready to use as LUNs. For help choosing other RAID schemes, see “Choosing RAID Schemes for LUNs” on page 30.

Step 6: Create a Metadata Array

Ten gigabytes (GB) of disk space is enough to store the metadata for a volume containing 10 million files, so a two-drive RAID 1 (mirrored) array is generally large enough to store the metadata for your SAN volume. If you dedicate a spare drive to this array to guarantee availability, then three drives are adequate for your SAN metadata. If all of your RAID arrays consist of four or more drives, you can convert one of these existing arrays into a small metadata array so you can reuse the extra drives.

- 1 If you don't have three spare drives or if all of the drives in your RAID systems already belong to RAID arrays, use the management application for your RAID system to delete an existing array.
- 2 Use two of the drives to create a new two-drive RAID 1 (mirrored) array.
- 3 Assign a third drive as a dedicated spare for the array.
- 4 You can use any leftover drives from the original array to create a separate array, or save them for use as spares.

For best performance, a RAID array used for metadata shouldn't be on the same RAID controller as an array used for user data.

Step 7: Install Xsan Software on Clients and Controllers

Take the *Xsan Install Disc* to each client computer connected to the SAN and install the Xsan software. Then install the Xsan software on the servers that will be standby metadata controllers. Last, install the Xsan software on the server that will be the primary metadata controller.

Note: If you use Server Assistant to set up a server as metadata controller, it prompts you to insert the *Xsan Install Disc* when it's time to install the Xsan software. The last computer you install the Xsan software on should be the primary metadata controller.

To install the Xsan file system and Xsan Admin application:

- Insert the *Xsan Install Disc*, double-click the Install Xsan.mpkg icon, and then follow the onscreen instructions to install the Xsan software.

If you're setting up new metadata controller, or a controller on which you have just performed a clean installation of Leopard Server, the setup assistant prompts you to insert the *Xsan Install Disc* when it's time to install the Xsan software.

To install just the Xsan file system without Xsan Admin:

- 1 Insert the *Xsan Install Disc*, double-click the Install Xsan.mpkg icon, and then follow the onscreen instructions until you reach the Custom Install pane.

If you're setting up new metadata controller, or a controller on which you have just performed a clean installation of Leopard Server, the setup assistant prompts you to insert the *Xsan Install Disc* when it's time to install the Xsan software.

- 2 In the Custom Install pane, deselect Xsan Admin, click Continue, and follow the remaining onscreen instructions to install just the Xsan file system.

To install Xsan software on a computer that has no keyboard or display:

- 1 Insert the *Xsan Install Disc* in the computer that doesn't have a keyboard or display.
- 2 Log in to a computer that does have a keyboard and display.
- 3 In the Finder, choose Go > Connect to Server and type `vnc://address` in the Server Address field, replacing *address* with the IP address or DNS name of the target computer.
- 4 Click Connect and enter the name and password of an administrator account on the target computer.
- 5 In the screen sharing window, open the *Xsan Install Disc* icon, double-click the Xsan.mpkg icon, and follow the onscreen instructions.

If you want to install just the Xsan file system, deselect Xsan Admin in the Custom Install pane.

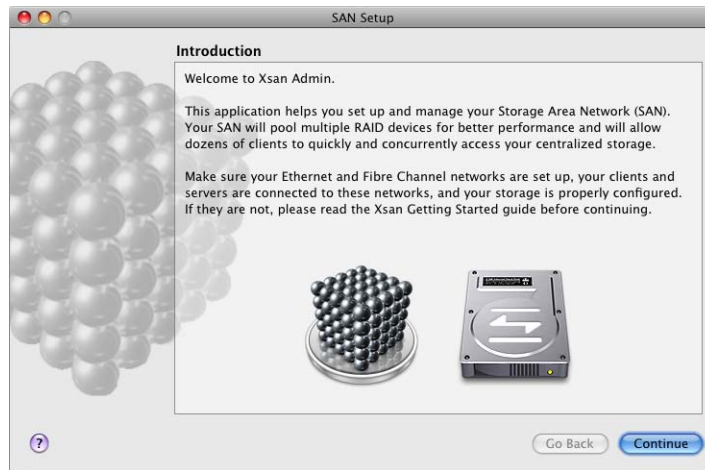
You can also use the Apple Remote Desktop application (which you can purchase separately) to install the Xsan software on one or more remote computers at once. For information about Apple Remote Desktop, go to www.apple.com/remotedesktop.

From the Command Line

For information about installing the Xsan software by using command-line tools in Terminal, see "Installing Xsan from the Command Line" on page 120.

Step 8: Configure the SAN

You use Xsan Admin to configure the SAN. The first time you use Xsan Admin on a computer, it automatically starts the Xsan setup assistant.



- 1 Open Xsan Admin (in /Applications/Server/).
Xsan Admin opens automatically when you finish installing the Xsan software on your primary metadata controller computer.
You can also open Xsan Admin on any computer with an intranet or Internet connection to your SAN computers. (You can use Xsan Admin on a computer that isn't connected to the SAN's private metadata network or its Fibre Channel network.)
- 2 In the Introduction pane, click Continue.
- 3 In the Initial SAN Setup pane, select "Configure new SAN".
For information about connecting to an existing SAN, see "Managing Multiple SANs" on page 52.
- 4 In the SAN Settings pane, type a name for the SAN, and then enter the SAN administrator's name and email address.
- 5 In the Add Computers pane, make sure all the computers that you want to include in the SAN are selected.
If a computer you want to include isn't listed, make sure you have installed Xsan on that computer, check that it is connected to the private metadata Ethernet network, and check the network settings in the computer's Network preferences. You can also try clicking Add Remote Computer and entering the computer's address.
- 6 In the Authenticate SAN Computers pane, choose how you will provide authentication information for the SAN computers:

- **Use same authentication information for all SAN computers:** Select this option to have Xsan Admin authenticate to all computers using the administrator name and address you enter in this pane.
 - **Authenticate to SAN computers one by one:** Select this option to authenticate to each computer individually.
- 7 In the Serial Numbers pane, enter your Xsan serial numbers.

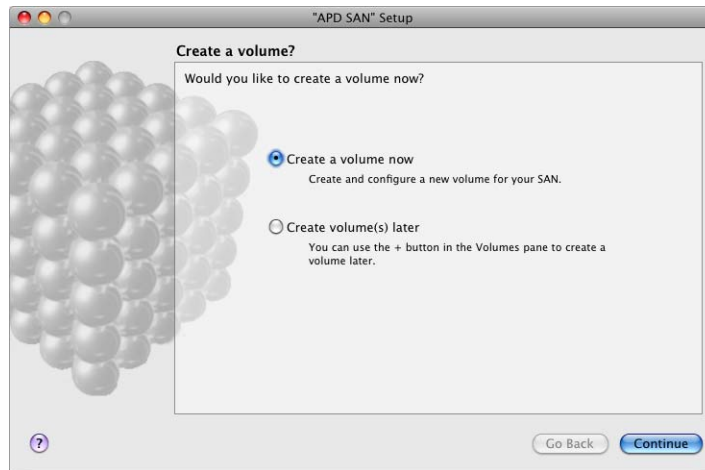
You can click Add Serial Number and type a number, or drag a text file containing serial numbers to the list.
 - 8 In the Choose Metadata Controllers pane, select just your primary and standby controllers. Deselect any client-only computers that appear in the list.
 - 9 If the Private Metadata Network pane appears, you can choose to have Xsan Admin manage the private network addresses for all SAN computers.
 - 10 If the SAN Network pane appears, choose your private network from the Metadata Network pop-up menu.

The SAN Network pane doesn't appear if in the previous step you chose to have Xsan Admin manage the private network addresses.
 - 11 Review the Summary pane, and if all settings are correct, click Continue.

To change a setting, click Go Back until you reach the pane where you can correct the setting. Then click Continue until you return to the Summary pane.

Step 9: Create a Volume

When the Xsan setup assistant finishes basic SAN configuration, it asks if you want to set up a volume.



- 1 In the "Create Volume" pane, select "Create a volume now" and click Continue.

If you want to set up volumes later, follow the instructions under "Adding a Volume to a SAN" on page 58.

- 2 In the "Volume Name and Type" pane, type a name for the volume and choose a volume type that matches the type of work the volume will support. Optionally, click Advanced Settings.

For the volume name, use only uppercase letters (A–Z), lowercase letters (a–z), numbers (0–9), and underscores (_). Don't include spaces or hyphens. The maximum length is 70 characters.

The volume type you choose determines how the setup assistant will configure affinity tags and storage pools on the volume. For information, see "How Xsan Storage Is Organized" on page 16.

If you click Advanced Settings, you can adjust the following volume settings.

- **Block Allocation Size:** If you're not sure what value to use, use the preset size or see "Setting the Block Allocation Size" on page 62.
- **Allocation Strategy:** Choose how storage for files is allocated among the storage pools that belong to the volume. If you choose Round Robin, each new request for space is assigned to the next available storage pool in turn. If you choose Fill, all space is allocated on the first storage pool until it is full, then on the second storage pool, and so on. If you choose Balance, space is allocated on the storage pool that has the most free space. For more information, see "Choosing an Allocation Strategy" on page 34.

- **Spotlight:** Enable this if you want Macintosh clients to be able to search the contents of the volume using Spotlight.
 - **Access Control Lists:** Leave this enabled if you want to use access control lists in Xsan Admin to control access to the volume and its contents.
 - **Windows ID Mapping:** If you will have Windows clients on your SAN, choose how they map user and group information to the Xsan-compatible user IDs and group IDs they need in order to access this volume. For more information, see “Mapping Windows User and Group IDs” on page 84.
- 3 In the Configure Volume Affinities pane (or the Configure Volume Storage pane, if you’re configuring a custom volume type), drag LUNs from the left column to the corresponding affinity tag (or custom storage pool) in the right column.
- a Drag the special metadata LUN you created (in Step 6, “Create a Metadata Array”) to the MetadataAndJournal affinity tag (or custom storage pool).
 - b Drag your other LUNs to the other affinity tags (or storage pools). To avoid wasting storage, all LUNs assigned to an affinity tag (or storage pool) should be the same size.
 - c When you finish, click Continue.

For information about the optimal number of LUNs to assign to an affinity tag, see “Assigning LUNs to Affinity Tags” on page 32.

You can select an affinity tag and click Settings to change the affinity tag name or other settings listed below. If you’re creating a custom volume, you can select it and click Storage Pool Settings to change the storage pool name or other settings listed below.

- **Affinity Tag (or Storage Pool Name):** Type the name for the affinity tag (or custom storage pool). If the OK button is disabled when you finish typing the name, the name is reserved; try another. For a list of reserved names, see Chapter 8.
 - **Use for:** Choose the types of data that can be stored on storage pools that have this affinity tag (or that can be stored on the custom storage pool).
 - **Stripe Breadth:** Specify how much data is written to or read from each LUN in storage pools that have this affinity tag (or each LUN in the custom storage pool) before moving to the next LUN. This value can affect performance. If you’re not sure what value to use, use the preset value or see “Setting the Storage Pool Stripe Breadth” on page 66.
- 4 If the Volume Failover Priority pane appears, arrange the list so as few SAN volumes as possible have the same metadata controller first on their failover priority lists, and then click Continue.
- 5 In the Setup Complete pane, click Continue.

Xsan Admin displays a summary of your SAN configuration, and the new volume is mounted and ready to use in the Finder on all clients and metadata controllers.

For information about creating additional volumes, see “Adding a Volume to a SAN” on page 58.

Step 10: (Optional) Customize SAN Status Notifications

Xsan is set initially to notify the administrator by email when the status of the SAN changes. You can have notifications sent to additional email addresses or text-messaging addresses, and you can choose conditions that trigger notification to each address. If you don’t want to customize notifications now, you can do it later. For instructions, see “Setting Up Status Notifications” on page 105.

Step 11: (Optional) Assign Folders to Affinity Tags

If you want to force files to be stored in a specific storage pools, assign the affinity tag of the pools to a folder. Then, files that users put in the folder are stored only on storage pools that have that affinity tag. For instructions, see “Setting Up a Folder Affinity” on page 59.

Step 12: (Optional) Set User and Group Quotas

You can set up quotas to control how much space on each SAN volume is used by each user or group. For instructions, see “Setting SAN User and Group Quotas” on page 85.

Setting Up an Xsan Administrator Computer

You can install the Xsan Admin application on any computer with Mac OS X v10.5 Leopard or Mac OS X Server v10.5 Leopard, and then use it to administer an Xsan 2 SAN remotely. The administrator computer must be able to connect to all SAN computers via your public intranet or the Internet. The administrator computer doesn’t need to be connected to the SAN’s private metadata network or the SAN’s Fibre Channel network.

To install just the Xsan Admin application:

- Insert the *Xsan Install Disc*, open the Other Installs folder, double-click the XsanAdmin.mpkg icon, and follow the onscreen instructions.

For information about using Xsan Admin to manage a SAN remotely, see “Managing Multiple SANs” on page 52.

Renaming a SAN

The SAN name appears in the Overview pane of Xsan Admin. The SAN name is initially set when the SAN is set up. You can change this name using Xsan Admin.

To change the name of a SAN:

- 1 Open Xsan Admin (in /Applications/Server/) and click Overview.
- 2 Choose Edit SAN Properties from the Action (gear) pop-up menu.
- 3 Type a name in the SAN Name field and click Save.

Removing a SAN

Follow these steps to take a SAN out of service, removing its metadata controllers, clients, and volumes.

WARNING: Removing a SAN destroys its volumes. Data stored on the volumes is no longer available.

To remove a SAN:

- 1 If you want files located on the SAN volumes to be available after you remove the SAN, back up the files.
- 2 Open Xsan Admin (in /Applications/Server/) and click Overview.
- 3 Choose Destroy SAN from the Action (gear) pop-up menu.

Managing Multiple SANs

You can use Xsan Admin to manage more than one Xsan 2 SAN. The computer with Xsan Admin doesn't need to be connected to the SAN's private metadata network or its Fibre Channel network, but must be able to connect to the SAN computers via your public intranet or the Internet.

To manage another SAN:

- 1 Open Xsan Admin and choose File > New.
- 2 Click Continue in the Introduction pane.
- 3 Select "Connect to existing SAN," click Continue, and follow the onscreen instructions.

Setting Up Additional SANs

You can use Xsan Admin to set up more than one SAN.

To add a new SAN:

- 1 Install the hardware, connect the Ethernet and Fibre Channel networks, set up the client computers, set up standby metadata controllers if you have them, set up the RAID systems, create a metadata array, and set up the primary metadata controller as instructed earlier in this chapter.

If you are setting up a SAN for the first time, see the planning guidelines in Chapter 2 and the detailed instructions at the beginning of this chapter.

- 2 Open Xsan Admin on the computer that you want to use to set up and manage the new SAN.
- 3 Choose File > New and follow the instructions in "Configure the SAN" on page 47.

You can use Xsan Admin and related command-line tools to expand, add, modify, check, and repair SAN volumes.

This chapter shows how you can expand an existing Xsan volume to provide more free space. It also contains information on volume and storage pool settings, and shows how to check and resolve volume integrity and fragmentation problems.

Adding Storage

To increase the storage on your SAN, you can:

- Add LUNs to existing storage pools
- Add storage pools to existing volumes
- Add new volumes

Adding LUNs to a storage pool increases the size of an existing volume, and can improve performance by increasing the number of RAID controllers and data paths between clients and storage. During the expansion, Xsan Admin stops the volume and unmounts it from clients while it adds the new LUNs.

Adding an entire storage pool to a volume also increases available storage, and also requires Xsan Admin to stop the volume and unmount it from clients.

If you create a volume based on a custom volume type, you work directly with storage pools. However, if you create a volume using a built-in Xsan volume type (for example, General File Server or Calendar Server Cluster), you don't deal directly with storage pools. Instead, you work with affinity tags that represent one or more storage pools. Xsan Admin organizes available LUNs into storage pools for you based on the performance requirements of the chosen volume type.

Preparing LUNs

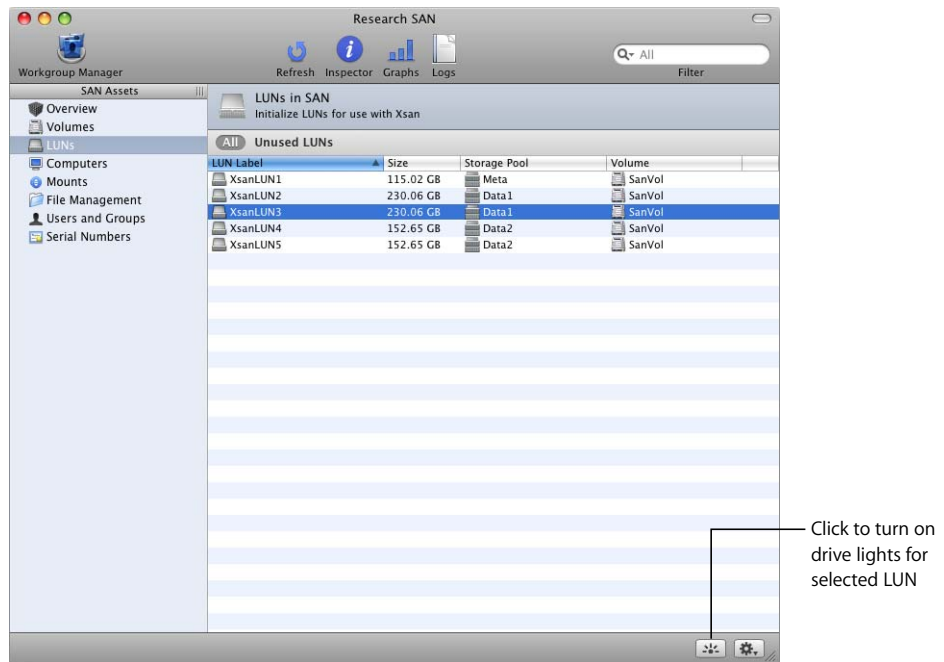
Each LUN in an Xsan volume is a RAID array. The way you set up your arrays depends on the storage device you are using.

If you're adding new RAID systems to your SAN, they may be ready to use right out of the box. Many RAID systems ship already configured as one or more RAID arrays, and each array should appear in the Xsan setup assistant as a usable LUN.

To create a different set of LUNs for your SAN, use the application that comes with your RAID system to create, for example, LUNs based on different RAID schemes or LUNs based on array stripes.

Finding the Drive Modules That Belong to a LUN

To see which drive modules belong to a LUN, you can use Xsan Admin to turn on the drive activity lights on the RAID system that hosts the LUN.



To find a LUN's drives:

- In Xsan Admin, select LUNs in the SAN Assets list, select a LUN in the list of LUNs, and click the "Identify LUN Using RAID Lights" button in the lower-right corner of the window.

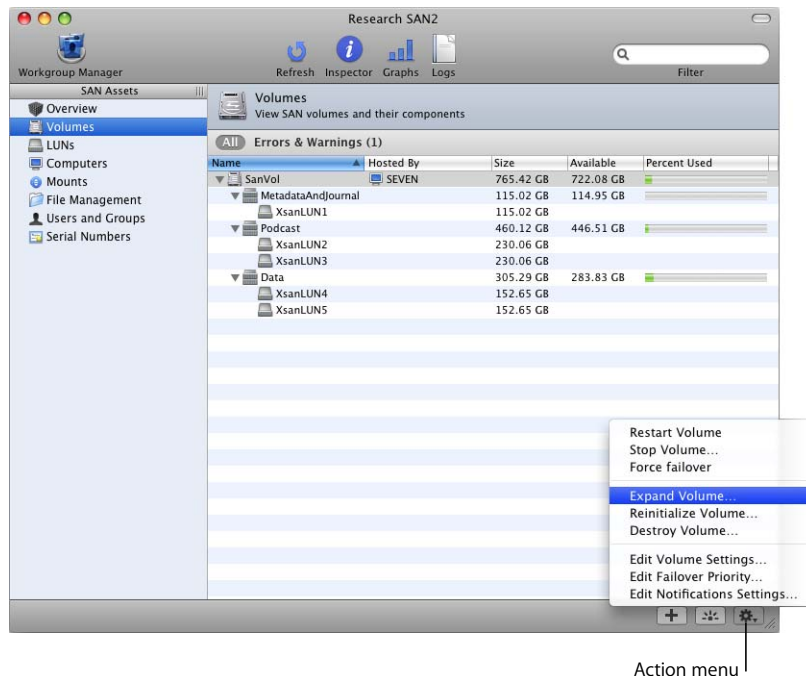
Adding LUNs to a Storage Pool

You can increase the size of a SAN volume by adding LUNs (RAID arrays or array slices) to storage pools in the volume. If you're expanding a volume that is based on a built-in volume type, you add LUNs to affinity tags and Xsan Admin assigns the LUNs to underlying storage pools for you. If you're expanding a custom volume, you add LUNs directly to storage pools.

Note: A storage pool can't contain more than 32 LUNs, the total number of LUNs in a volume can't be greater than 512, and you can't add LUNs to a storage pool that contains only journal data or metadata.

Choosing Compatible LUNs

LUNs you add to an existing storage pool must be at least as large as the LUNs already in the pool. On the other hand, if a new LUN is larger than the other LUNs in the pool, its extra capacity can't be used. So, always try to add LUNs that are identical or similar in performance and capacity to the LUNs already in the storage pool. Mixing LUNs of different sizes or speeds in the same storage pool wastes capacity and can degrade performance.



To add a LUN to a storage pool:

- 1 If you haven't already done so, connect the RAID system that hosts the LUN to the SAN Fibre Channel network and power on the device.
- 2 In Xsan Admin, select Volumes in the SAN Assets list.
- 3 Select the volume in the list and choose Expand Volume from the Action (gear) pop-up menu.
- 4 In the Label LUNs pane of the assistant, choose whether you want to label unlabeled LUNs individually or sequentially based on a label prefix.

If you use a label prefix, Xsan Admin adds a number to the end of the prefix to create a label for each LUN. For example, if you use the prefix "LUN," your LUNs will be labeled "LUN1," "LUN2," etc. Already labeled LUNs are not affected.

If you choose to label LUNs individually, click Edit LUN Label on the next pane and type a new label.

If your LUNs are already labeled, they are not changed.
- 5 In the Configure Volume Storage pane, drag the new LUNs to affinity tags (or storage pools, if this is a custom volume).
- 6 Click Continue to unmount and stop the volume, add the new storage, and remount the expanded volume.

From the Command Line

You can also add LUNs to a storage pool by modifying the associated volume configuration file and using the `cvlabel` command in Terminal. For more information, see the `cvfs_config` and `cvlabel` man pages or "Xsan Configuration Files" on page 131 and "Labeling, Listing, and Unlabeling LUNs (cvlabel)" on page 126.

Rearranging Fibre Channel Connections

If you need to rearrange your Fibre Channel connections when adding LUNs to your SAN, always unmount SAN volumes from clients *before* you disconnect Fibre Channel cables or turn off Fibre Channel switches. Otherwise, if you unplug or interrupt a Fibre Channel connection between a client and a mounted volume, you might cause problems with client applications that are using the volume and make the volume difficult to remount.

Adding a Storage Pool to a Volume

You can add free space to a SAN volume by adding a storage pool to the volume.

If you're expanding a volume that is based on a built-in volume type, you actually create a new affinity tag and add LUNs to it. Xsan Admin creates and organizes new storage pools within that tag for you.

If you're expanding a custom volume, you create new storage pools directly and add LUNs directly to them.

To add a storage pool:

- 1 If necessary, connect the RAID systems that host the storage pool's LUNs to the SAN Fibre Channel network and power on the device.
- 2 In Xsan Admin, select Volumes in the SAN Assets list.
- 3 Select the volume in the list and choose Expand Volume from the Action (gear) pop-up menu.
- 4 In the Label LUNs pane of the assistant, choose whether you want to label any unlabeled LUNs individually, or sequentially based on a label prefix.

If you use a label prefix, Xsan Admin adds a number to the end of the prefix to create a label for each LUN. For example, if you use the prefix "LUN," your LUNs will be labeled "LUN1," "LUN2," etc. Already-labeled LUNs aren't affected.

If you choose to label LUNs individually, click Edit LUN Label on the next pane and type a new label.

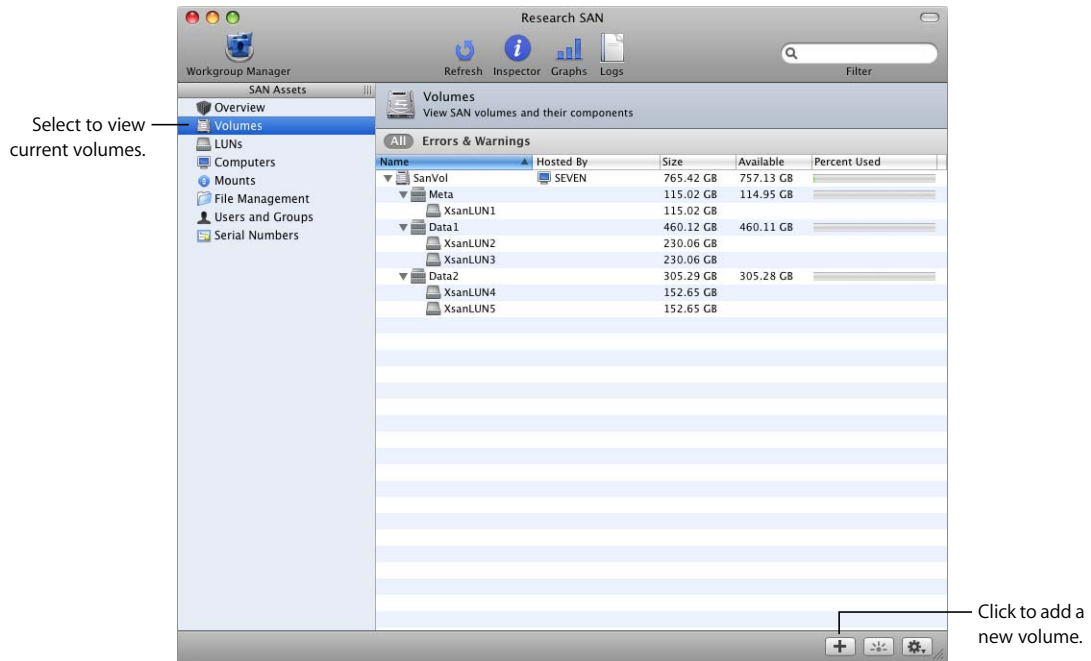
- 5 In the Configure Volume Storage pane, click New Affinity Tag (or New Pool) to add an affinity tag (or storage pool), and then drag LUNs to the new tag (or pool).
- 6 Click Continue to unmount and stop the volume, add the new storage, and remount the expanded volume.

From the Command Line

You can also add a storage pool by modifying the associated volume configuration file in Terminal. For more information, see the `cvfs_config` man page or "Xsan Configuration Files" on page 131.

Adding a Volume to a SAN

A single Xsan SAN can provide access to multiple volumes.



To add a volume:

- 1 In Xsan Admin, select Volumes in the SAN Assets list and click the Add Volume (+) button.
- 2 In the Volume Name and Type pane of the assistant, type a name for the volume and choose a volume type that matches the kind of work the volume will support. Xsan Admin sets the underlying volume settings accordingly.
- 3 In the Configure Volume Affinities pane (Configure Volume Storage if you chose the custom volume type), drag LUNs to affinity tags (or storage pools).
- 4 On the Volume Failover Priority pane, drag the controller that you want to host the volume whenever possible to the top of the list. Arrange the other controllers in descending order.

For information on advanced settings, see “Changing Advanced Volume Settings” on page 61 or “Changing Storage Pool Settings” on page 65.

When you’re finished, the volume is automatically mounted on all SAN clients.

From the Command Line

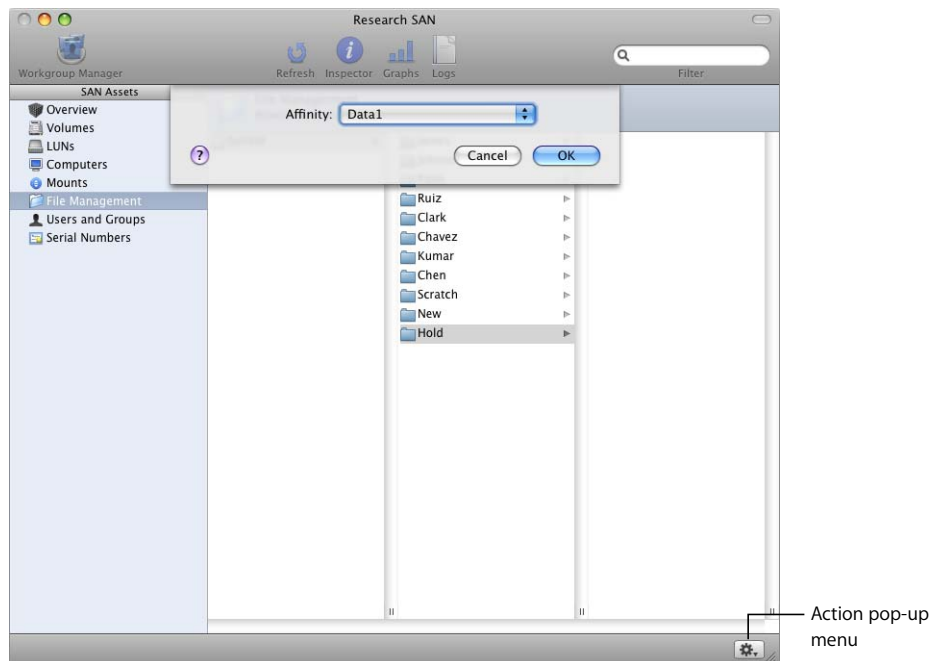
You can also add a volume by setting up a configuration file. For more information, see the `cvfs_config` man page or “Xsan Configuration Files” on page 131.

Setting Up a Folder Affinity

Every storage pool in a volume has an affinity tag. You can use the tag to be sure that files in a folder are stored on a particular storage pool. Files folders without affinities are stored in the next available storage pool according to the volume's allocation strategy.

You might have some storage pools that are larger, faster, or better protected than others. Using affinities, you can make sure that an application or task that needs speed or extra protection stores its files on a suitable storage pool.

Using Xsan Admin, you can choose an affinity for an existing folder or create a new folder with an affinity.



To assign an affinity tag to a folder:

- 1 In Xsan Admin, select File Management in the SAN Assets list.
- 2 Select the folder in the columns that list the volume's contents, choose Set Affinity from the Action (gear) pop-up menu, and choose an affinity tag.

If the folder doesn't already exist, you can choose New Folder from the Action (gear) pop-up menu, type a folder name, and choose an affinity tag.

From the Command Line

You can also assign an affinity to a folder using the `cvmkdir` command in Terminal. For more information, see the `cvmkdir` man page.

Changing a Folder's Storage Pool Affinity

You can use Xsan Admin to change a folder's affinity so all new files placed in the folder are stored on a new storage pool.

To change a folder affinity:

- 1 In Xsan Admin, select File Management in the SAN Assets list.
- 2 Select the folder, choose Set Affinity from the Action (gear) pop-up menu, and then choose the new affinity tag.
- 3 Click OK.

Files already in the folder are not automatically moved to the new storage pool. To move the files that were in the folder to a storage pool with the new affinity tag, use the `snfsdefrag` command in Terminal. For information and an example, see the `snfsdefrag` man page.

Removing an Affinity

You can undo a folder's affinity for a storage pool by choosing None for the folder affinity.

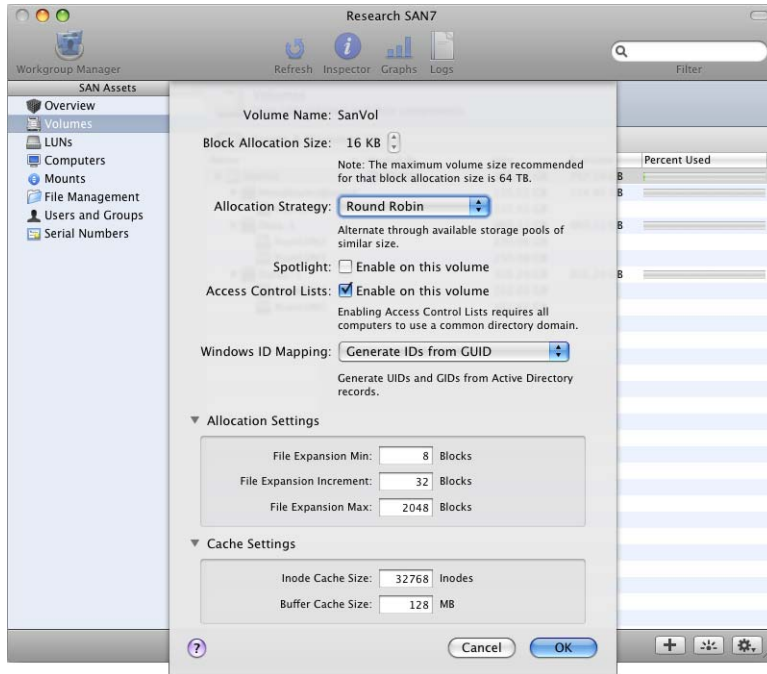
To remove the affinity tag from a folder:

- 1 In Xsan Admin, select File Management in the SAN Assets list.
- 2 Select the folder, choose Set Affinity from the Action (gear) pop-up menu, and choose None for the affinity tag.

Changing Advanced Volume Settings

If your SAN volume has special configuration requirements, you can change the standard volume settings when you create a volume. You can also change these settings for an existing volume, with the exception of the volume name and block allocation size.

Note: To change a volume's name or block allocation size, you must destroy and recreate the volume.



To view or change volume settings:

- Select Volume in the SAN Assets list, select the volume in the list, and choose Edit Volume Settings from the Action (gear) pop-up menu.

The following sections contain information about each setting.

Setting the Block Allocation Size

Xsan uses the volume block allocation size together with the storage pool stripe breadth to decide how to write data to a volume. If you create a volume based on a built-in volume type (for example, General File Server or Calendar Server Cluster), Xsan Admin sets optimal values for you. For most volumes, the preset block allocation size and storage pool stripe breadth result in good performance. However, in some cases you might be able to improve read or write performance by adjusting these settings to suit a specific application. For example, if your application reads and writes small blocks of data, you might improve performance by choosing a correspondingly small block allocation size.

For example, if your application reads and writes 16 KB blocks of data, you can try adjusting the block allocation size to 16 KB and then calculate a corresponding stripe breadth for the volume's storage pools using this formula:

$$\text{stripe breadth} = (\text{transfer size} / \text{number of LUNs}) / \text{block allocation size}$$

where stripe breadth is expressed in blocks, and transfer size and block allocation size are expressed in bytes.

Computers running Mac OS X or Mac OS X Server have an optimal transfer size of 1 MB. If in the example there are four LUNs in the volume's data storage pools, the calculation for becomes:

$$\text{stripe breadth} = (1048576 / 4) / 16384$$

This gives a stripe breadth of 16.

Changing the Volume Allocation Strategy

You can change the allocation strategy for a volume to choose how storage for new files or additional storage for existing files is allocated on the storage pools that belong to the volume.

- **Round Robin:** Each new request for space is assigned to the next available storage pool in the volume.
- **Fill:** All data is stored on the first storage pool until it is full, then on the next storage pool, and so on.
- **Balance:** New data is written to the storage pool that has the most free space.

Enabling or Disabling Spotlight on a Volume

You can use Xsan Admin to control whether a volume is indexed and searchable using Spotlight.

To enable or disable Spotlight on a volume:

- 1 In Xsan Admin, select Volumes in the SAN Assets list.
- 2 Select the volume and choose Edit Volume Settings from the Action (gear) pop-up menu.
- 3 Set or clear the checkbox next to Spotlight and click OK.

Enabling or Disabling Access Control Lists

You can use Xsan Admin to specify whether the Xsan file system uses access control lists (ACLs) on a volume.

Xsan 2 clients, Xsan 1.4 clients, and Windows StorNext clients recognize ACLs. UNIX clients ignore ACLs on Xsan volumes. If you have a mix of Windows clients and Xsan clients, they must all be bound to the same directory domain, whether provided by Open Directory configured as a primary domain controller (PDC) or by Windows Active Directory.

Note: If you enable ACLs but your SAN includes clients that don't support them, don't use those clients to change file or folder ownership information or inconsistencies might result.

To enable or disable ACLs:

- 1 In Xsan Admin, select Volumes in the SAN Assets list.
- 2 Select the volume and choose Edit Volume Settings from the Action (gear) pop-up menu.
- 3 Set or clear the checkbox next to Access Control Lists and click OK.

Changing the Windows ID Mapping

If you have Windows clients on your SAN, this setting determines how they map user and group information to the Xsan-compatible user IDs and group IDs they need in order to access this volume. For more information, see “Mapping Windows User and Group IDs” on page 84.

Changing Advanced Allocation and Cache Settings

Volume settings that control the allocation of space for growing files and the caching of file-related data structures are set by Xsan Admin to suit the type of volume you set up. If necessary, you can use Xsan Admin to adjust these advanced allocation and caching settings for a volume.

Important: Do not adjust these settings unless you understand their role in volume performance or you are directed to change them by Apple support personnel.

To change advanced volume settings:

- In Xsan Admin, select Volumes in the SAN Assets list, select the volume, and choose Edit Volume Settings from the Action (gear) pop-up menu.

File Expansion Min: The number of storage blocks added to the file for the first expansion request.

File Expansion Increment: The number of storage blocks by which the expansion request is increased for each subsequent request.

File Expansion Max: The maximum expansion request that is allowed.

Inode Cache Size: The maximum number of inode data structures that can be cached on the volume by the metadata controller.

Buffer Cache Size: The amount of memory that the metadata controller can allocate for storing a volume’s metadata.

Renaming a Volume

To change the name of a volume, you need to reinitialize the volume with the new name. You can't rename an Xsan volume using the Finder.

WARNING: Reinitializing a volume destroys all information stored on the volume.

To rename a volume:

- In Xsan Admin, select Volumes in the SAN Assets list, select the Volume, and choose Reinitialize Volume from the Action (gear) pop-up menu.

Changing Storage Pool Settings

The SAN Setup assistant chooses appropriate storage pool settings based on the type of volume you create. To specify storage pool settings yourself, choose an affinity tag or storage pool on the Configure Volume Affinities pane of the SAN Setup assistant when you create the volume and click the Settings button below the list.

The best way to set up a SAN is to plan its organization carefully before you set it up, including settings for the storage pools that make up its volumes. To change storage pool settings for an existing volume; you must destroy and recreate the volume.

Choosing the Data Type for an Affinity Tag

The “Use for” setting specifies what type of data can be stored on a particular storage pool or on any storage pool with the corresponding affinity tag.

To set a storage pool's data type:

- The data type must be set when the storage pool (or affinity tag) is created; it can't be changed for an existing pool. To set the data type, click the Settings (or Pool Settings) button on the Configure Volume Affinities pane of the setup assistant when you first add the pool. Type a name for the affinity tag (or custom storage pool). If the OK button is disabled when you finish typing the name, the name is reserved; try another. For a list of reserved names, see Chapter 8.

Setting the Storage Pool Stripe Breadth

Xsan uses both the storage pool stripe breadth and the volume block allocation size to decide how to write data to a volume. For most SANs, the default values for storage pool stripe breadth and volume block allocation size result in good performance. However, in some cases you might be able to improve read and write performance by adjusting these values to suit a specific application.

The stripe breadth of a storage pool is the number of file allocation blocks that are written to a LUN in the pool before moving to the next LUN. To choose an efficient stripe breadth, you need to consider three factors:

- The Mac OS X and Mac OS X Server 1 MB data transfer size
- The number of LUNs in the storage pool
- The size of the data blocks written and read by the critical application that uses the volume (as reflected in the block allocation size for the volume)

Knowing these values, choose a stripe breadth using this formula:

$$\text{stripe breadth} = (\text{transfer size} / \text{number of LUNs}) / \text{block allocation size}$$

where stripe breadth is expressed in blocks, and transfer size and block allocation size are expressed in bytes.

For example, if you use an application such as Final Cut Pro to move large amounts of video data to and from a storage pool consisting of 4 LUNs, choose a block allocation size 256 KB and use the equation to find the stripe breadth of 1 block. Then, Xsan writes 256 KB of data, in turn, to each LUN in your video storage pool.

To set a storage pool's stripe breadth:

- The stripe breadth must be set when the volume is created; it can't be changed for an existing volume. To set the stripe breadth, click the Settings button below the Affinity Tag list on the Configure Volume Affinities pane of the setup assistant when you create the volume.

Checking Volume Fragmentation

When you create a file, Xsan divides the file into pieces and distributes these pieces efficiently over the LUNs that make up one of the volume's storage pools. Over time, as the file is modified, its pieces become scattered in less efficient arrangements. You can use the `cvfsck` utility to check the state of file fragmentation on your volumes.

To check volume fragmentation:

- 1 Open Terminal (in /Applications/Utilities/).
- 2 If you are not working at the SAN controller computer, use SSH to log in to the controller remotely:

```
$ ssh user@computer
```

where *user* is an administrator user on the controller computer and *computer* is the controller's name or IP address.
- 3 Run the `cvfsck` command-line utility with the `-f` option:

```
$ sudo cvfsck -f volume
```

For more information, see the `cvfsck` man page.

Defragmenting a Volume

Defragmenting a file reassembles its pieces into the most efficient arrangement. You can use the `snfsdefrag` command to defragment a file, a folder, or an entire volume.

To defragment a file, folder, or volume:

- 1 Open Terminal (in /Applications/Utilities/).
- 2 If you are not working at the SAN controller computer, use SSH to log in to the controller remotely:

```
$ ssh user@computer
```

where *user* is an administrator user on the controller computer and *computer* is the controller's name or IP address.
- 3 Run the `snfsdefrag` command.

To defragment one or more individual files:

```
$ sudo snfsdefrag filename [filename ... ]
```

To defragment an entire folder:

```
$ sudo snfsdefrag -r folder
```

To defragment an entire volume, set *folder* to the volume name.

For more information see the `snfsdefrag` man page or “Defragmenting a File, Folder, or Volume (snfsdefrag)” on page 128.

Checking the Integrity of a Volume

If SAN users are having trouble accessing files, you can use the `cvfsck` command to check the integrity of a volume, its metadata, and its files.

To check a volume:

- 1 Open Terminal (in /Applications/Utilities/).
- 2 If you are not working at the SAN controller computer, use SSH to log in to the controller remotely:

```
$ ssh user@computer
```

where *user* is an administrator user on the controller computer and *computer* is the controller's name or IP address.
- 3 Run the `cvfsck` command-line utility (in /Library/Filesystems/Xsan/bin/) to check the volume without making repairs:

```
$ sudo cvfsck -vn volume
```

You'll see a warning that the journal is active; this is normal.

For more information about this command, see the `cvfsck` man page.

Repairing a Volume

If the `cvfsck` utility reveals problems with a volume, you can use the same command to repair the volume.

To repair a volume:

- 1 Stop the volume.
Open Xsan Admin, select the volume, and click Stop Volume in the Action (gear) pop-up menu.
- 2 Open Terminal (in /Applications/Utilities/).
If you are not working at the SAN controller computer, use SSH to log in to the controller remotely:

```
$ ssh user@computer
```

where *user* is an administrator user on the controller computer and *computer* is the controller's name or IP address.
- 3 Run the `cvfsck` Command-line utility (in /Library/Filesystems/Xsan/bin/) to replay the events that are recorded in the file system journal:

```
$ sudo cvfsck -j volume
```

- 4 Check the volume to see if additional repairs are required:

```
$ sudo cvfsck -vn volume
```

- 5 If the report generated in the previous step lists additional problems, type the following command to perform a full check and repair of the volume:

```
$ sudo cvfsck -vw volume
```

For more information about this command, see the `cvfsck` man page.

Destroying a Volume

You can destroy a volume so you can reuse its LUNs to create new volumes.

To destroy a volume:

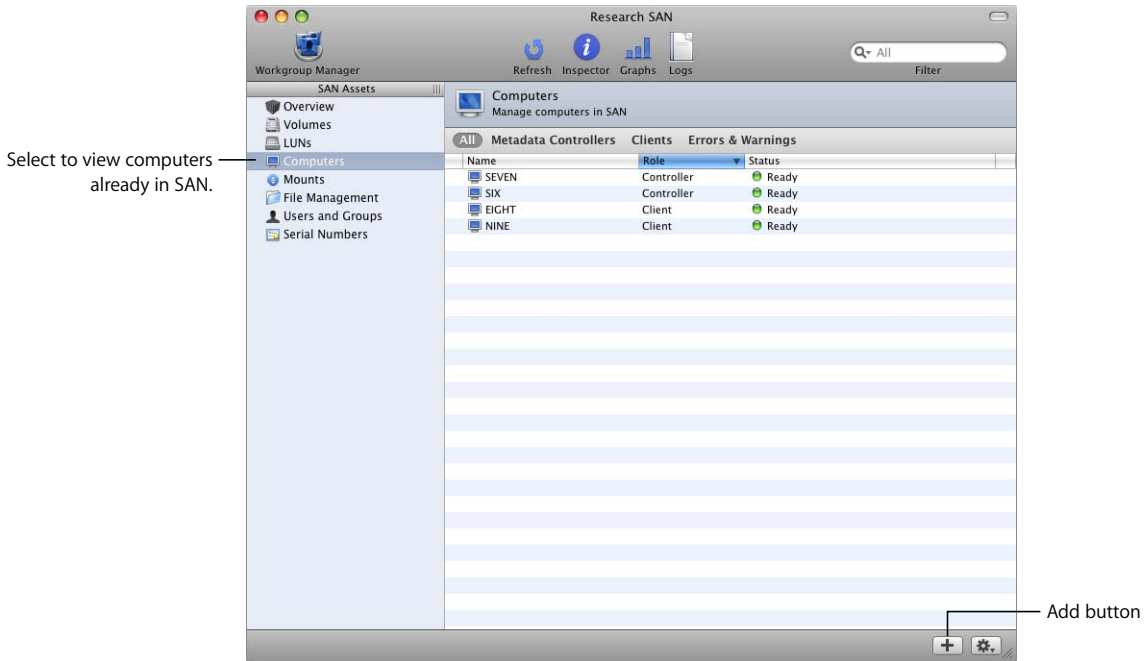
- In Xsan Admin, select Volumes in the SAN Assets list, select the volume in the list, and choose Destroy Volume from the Action (gear) pop-up menu.

You can use Xsan Admin and related command-line tools to add, control, and remove client computers and their users.

Xsan clients are computers that have Fibre Channel connections to a SAN. SAN users log in to client computers to access files stored on SAN volumes. This chapter shows you how to add clients, control client access to volumes, and manage user quotas.

Adding a Client

Before a computer can use a SAN volume, you need to add the computer to the SAN as a client. These instructions show you how to add a client computer to an existing SAN.

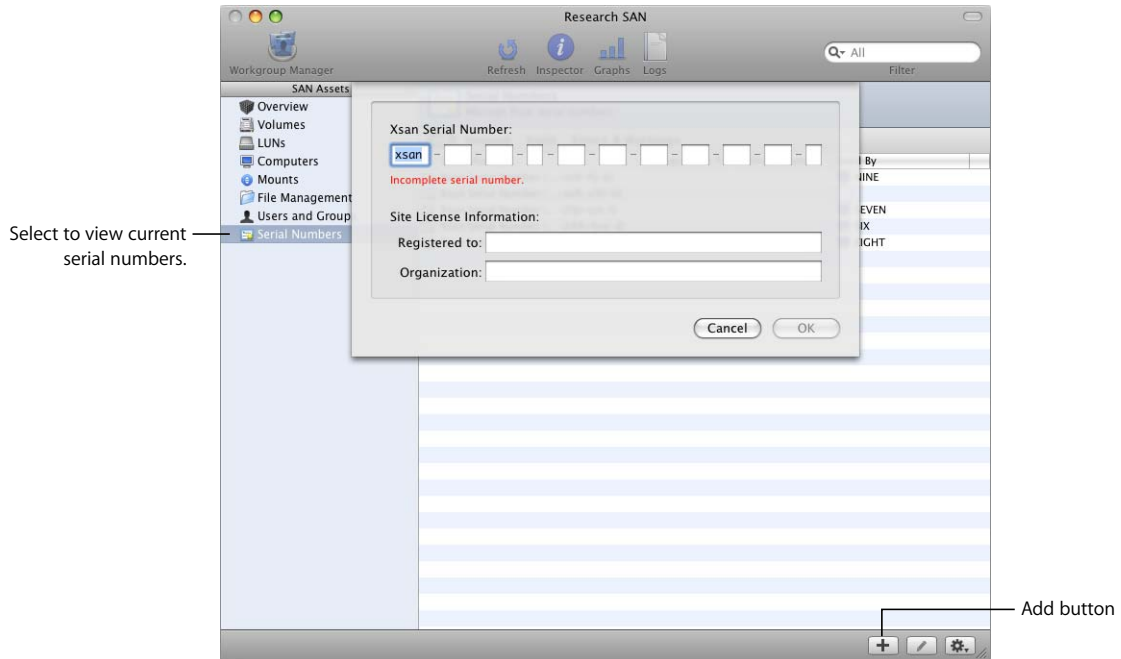


To add a client computer to a SAN:

- 1 Connect the client to the SAN's Fibre Channel and Ethernet networks.
- 2 Install the Xsan software on the client.
- 3 Open Xsan Admin, select Computers in the SAN Assets list, and click the Add (+) button.
- 4 In the Add Computers pane of the assistant, make sure there's a check next to the new client in the list, and then click Continue.
If the client isn't in the list, click Add Remote Computer and add it.
- 5 In the Authenticate Clients pane, enter the administrator name and password for the client and click Continue.
- 6 If there are no licenses available, the Licenses pane appears so you can add one.
- 7 In the Choose Metadata Controllers pane, make sure there's a no check next to the new client in the list, and then click Continue.
- 8 Click Continue on the Summary pane.

Adding an Xsan Serial Number

If you have purchased additional Xsan software serial numbers for client computers that you want to add to your SAN, you can enter the numbers in Xsan Admin.



To add an Xsan serial number:

- 1 In Xsan Admin, select Serial Numbers in the SAN Assets list.
- 2 Click the Add (+) button.
- 3 Type the serial number, registered owner, and organization information provided by Apple, and then click OK.

If you have serial numbers in a text file, you can simply drag the file to the Serial Number list in Xsan Admin.

Moving a Client to a Different SAN

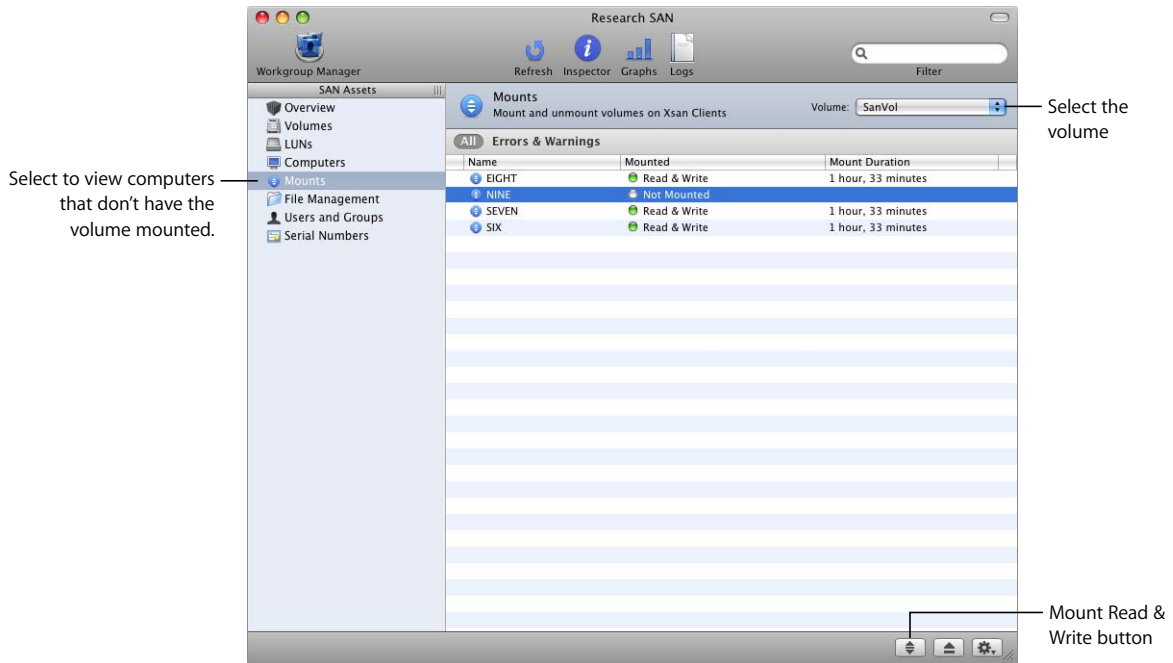
You can move a client from one Xsan SAN to a different SAN on the same Ethernet subnet and Fibre Channel network.

To move a client computer to a different SAN:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the computer in the list and choose “Remove computer from SAN” from the Action (gear) pop-up menu.
- 3 In Xsan Admin, open the window for the SAN you want to move the computer to.
- 4 In the new SAN window, select Computers in the SAN Assets list and click the Add (+) button.
- 5 Make sure there’s a check next to the computer in the list, and click Continue.
- 6 In the Authenticate Clients pane, enter the administrator name and password for the client and click Continue.

Mounting a Volume on a Client

Volumes are automatically mounted on SAN computers when you create the volumes. If, however, you explicitly unmount a volume from a client, you need to mount it again in order to restore access.



To mount an Xsan volume on a client:

- 1 In Xsan Admin, select Mounts in the SAN Assets list.
- 2 Select the client in the list.
- 3 Select the volume in the Volume pop-up menu.
- 4 To allow the client to modify files on the volume, click the Mount Read & Write button or choose Mount Read & Write from the Action (gear) pop-up menu.

To prevent the client from modifying the volume, Choose Mount Read Only from the Action (gear) pop-up menu.

A volume remains mounted until you unmount it or until the user on the client unmounts it, even if the user logs out or the client computer is restarted. If a user unmounts the volume, it is remounted automatically in a few moments.

From the Command Line

You can also mount a volume on a client using the `xsanctl` command in Terminal. For more information, see the `xsanctl` man page or “Mounting an Xsan Volume” on page 130.

Changing Mount Options

You can use Xsan Admin to adjust settings that can affect volume access performance.

To change mount options:

- 1 In Xsan Admin, select Mounts in the SAN Assets list.
- 2 Select the client in the list.
- 3 Select the volume in the Volume pop-up menu.
- 4 Choose Edit Mount Options from the Action (gear) pop-up menu.

Directory cache size: Controls the number of file system directory entries that are cached on the client for each SAN volume. Increasing this value makes sense if the volume contains large number of small files (for example, if the volume is hosting a home directory server, mail server, or calendar server).

Client worker threads: Controls the number of processes used to communicate with the volume. You might increase this if you are mounting many volumes on a client.

Delay access time updates until files are closed: Lets you increase performance by reducing the number of access time updates on a file that is read frequently (streaming video, for example). If not enabled, file access time is updated every time the file is read.

For more information, see the descriptions of these parameters in the `mount_acfs` man page.

Managing Users and Groups

Depending on how you first set up your SAN, you can manage users and groups using either Xsan Admin or Mac OS X Server's Workgroup Manager.

Managing Users and Groups with Xsan Admin

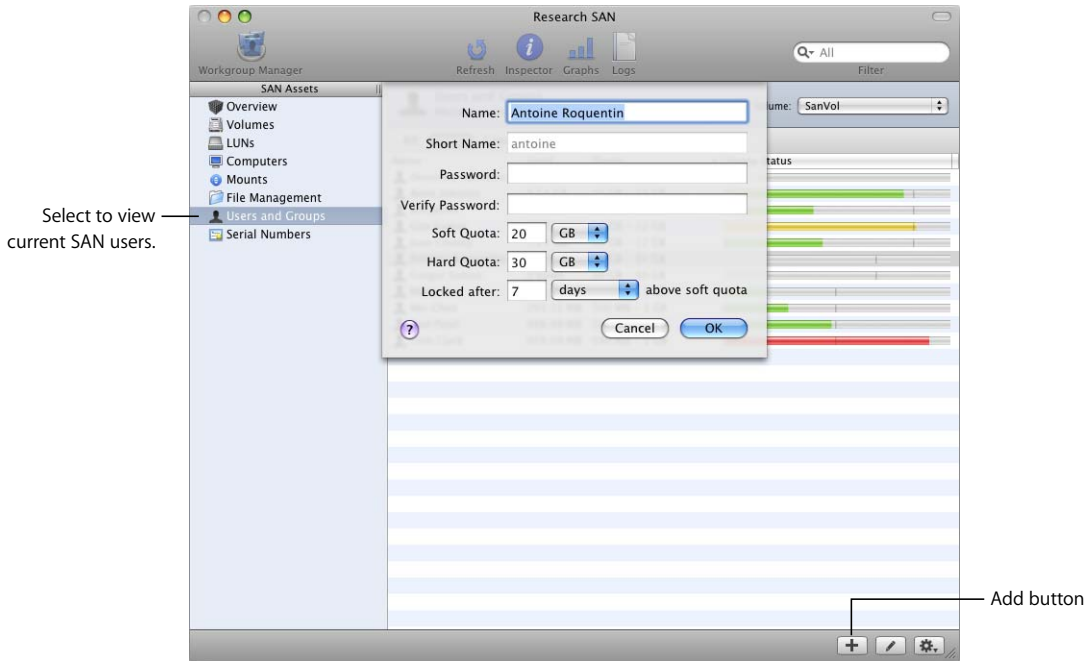
When you first set up your SAN controllers, you can use Xsan Admin to manage your SAN user and group accounts. An Open Directory master is created on your primary metadata controller, with replicas on your standby controllers. You can then use Xsan Admin to create and remove user and group accounts.

Managing Users and Groups with Workgroup Manager

If you choose not to manage accounts using Xsan Admin (for example, if you have an existing directory of user and group accounts), use the appropriate directory management tool to add or delete accounts. For example, if you have an existing Open Directory server, you can use Workgroup Manager in Mac OS X Server to manage user and group accounts and Directory Utility to bind SAN computers to that directory. See the Mac OS X Server documentation suite for details.

Adding SAN Users

Only users in the SAN's directory can log in to a client computer and access Xsan volumes. You can use Xsan Admin to add users to your SAN.



Note: These instructions apply only if, during initial SAN setup, you chose to use Xsan Admin to manage users and groups. If you have a different directory configuration, use the management software for your directory to add user accounts.

To add a user:

- 1 In Xsan Admin, select Users and Groups in the SAN Assets list.
Users and Groups appear only if you chose to have Xsan manage users and groups during initial setup.
- 2 Click the Users button above the list of users and groups.
- 3 Click the Add (+) button.
- 4 Enter a user name and password and click OK.

Deleting SAN Users

Only users in the SAN's directory can log in to a client computer and access Xsan volumes. You can use Xsan Admin to delete users from your SAN directory.

Note: These instructions apply only if, during initial SAN setup, you chose to use Xsan Admin to manage users and groups. If you have a different directory configuration, use the management software for your directory to delete user accounts.

To delete a user:

- 1 In Xsan Admin, select Users and Groups in the SAN Assets list.
- 2 Select the user in the list and choose Delete User or Group from the Action (gear) pop-up menu.

Creating Groups

You can create groups of users to simplify user management. If you chose to have Xsan manage your users and groups, you already have a group named Workgroup that contains all of your users.

Note: These instructions apply only if, during initial SAN setup, you chose to use Xsan Admin to manage users and groups. If you have a different directory configuration, use the management software for your directory to add group accounts.

To add a group:

- 1 In Xsan Admin, select Users and Groups in the SAN Assets list.
- 2 Click the Groups filter button above the list of users and groups.
- 3 Click the Add (+) button.
- 4 Enter a group name and password.
- 5 Select the checkbox next to the users who belong to the group.
- 6 Click OK.

Deleting Groups

You can use Xsan Admin to delete a group.

Note: These instructions apply only if, during initial SAN setup, you chose to use Xsan Admin to manage users and groups. If you have a different directory configuration, use the management software for your directory to delete group accounts.

To delete a group:

- 1 In Xsan Admin, select Users and Groups in the SAN Assets list.
- 2 Select the group in the list and choose Delete User or Group from the Action (gear) pop-up menu.

Changing Group Membership

You can use Xsan Admin to change the members of a group at any time.

Note: These instructions apply only if, during initial SAN setup, you chose to use Xsan Admin to manage users and groups. If you have a different directory configuration, use the management software for your directory to modify group membership.

To change a group's membership:

- 1 In Xsan Admin, select Users and Groups in the SAN Assets list.
- 2 Select the group in the list and click the Edit button in the lower-right corner of the window.
- 3 Select the checkbox next to a user to add the user to the group. Deselect the checkbox to remove the user.
- 4 Click OK.

Controlling Client and User Access

To control access to information on SAN volumes, you can:

- Use the Finder's Get Info window to apply basic access controls to a file or folder.
- Use Xsan Admin or Mac OS X Server's Server Admin application to apply a full set of access control list restrictions.
- Unmount a SAN volume from selected client computers (volume-level control).
- Restrict a client computer to read-only access to a volume (volume-level control).
- Remove a client from a SAN (SAN-level control).

Controlling File and Folder Access Using the Finder

To restrict access to a file or folder on an Xsan volume, you can use the Get Info window in the Finder.

To assign permissions using the Finder:

- In a Finder window, select the file or folder, choose File > Get Info, and look under Sharing & Permissions.

Controlling File and Folder Access Using Xsan Admin

To restrict user access to specific items on a SAN volume, you can use Xsan Admin to adjust permissions using access control lists (ACLs).

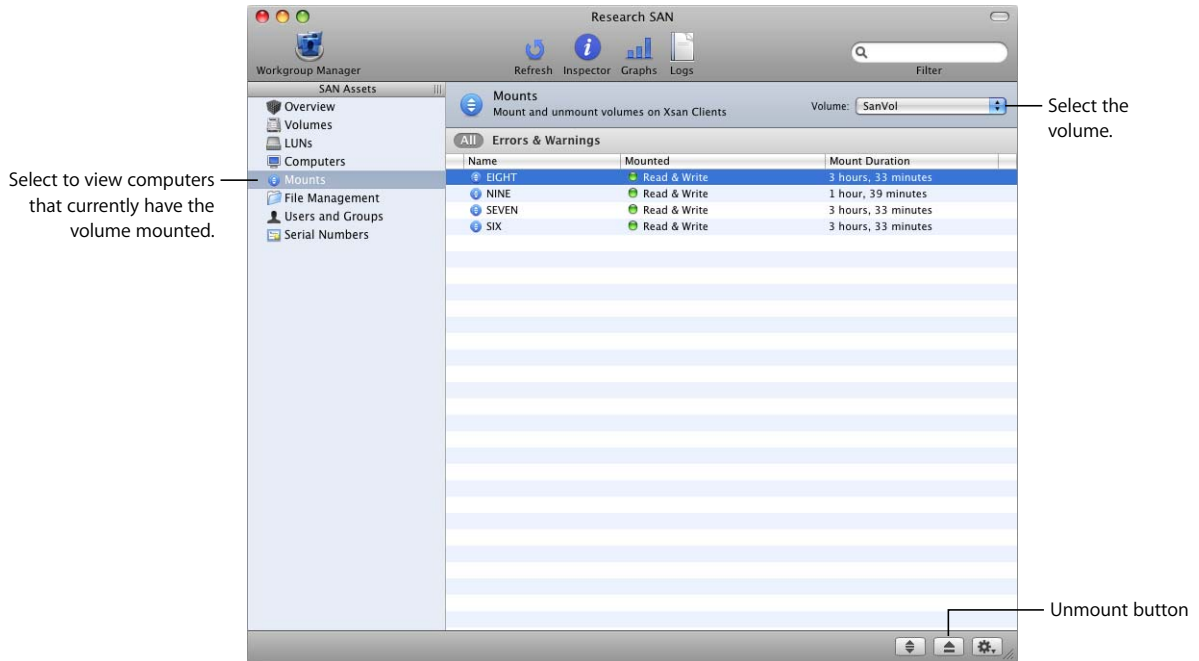
To assign permissions using Xsan Admin:

- 1 Make sure ACLs are enabled on the volume. For help, see "Enabling or Disabling Access Control Lists" on page 63.
- 2 In Xsan Admin, select File Management in the SAN Assets list.

- 3 Select the file or folder you want to protect and choose Set Permissions from the Action (gear) pop-up menu.

Unmounting a Volume on a Client

To prevent a client from accessing a volume, you can unmount the volume from the client. Clients can't mount SAN volumes themselves; only an administrator can mount a SAN volume on a client.



Note: A user can unmount a SAN volume from a client computer temporarily by ejecting it in the Finder like any other mounted volume. However, the volume is automatically remounted after a few moments. To be sure that a volume remains unmounted, use Xsan Admin to unmount it.

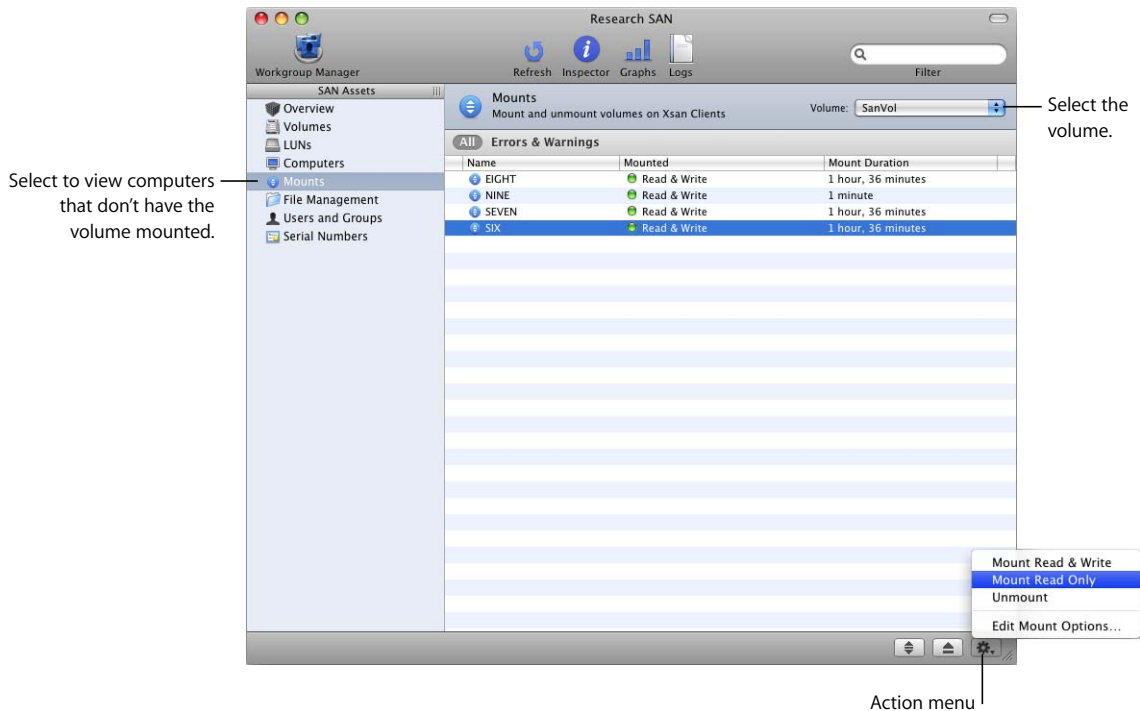
To unmount a volume:

- 1 In Xsan Admin, select Mounts in the SAN Assets list.
- 2 Choose the volume from the Volume pop-up menu.
- 3 Select the client in the list and click the Unmount button.

To select more than one client to unmount, hold down the Command or Shift key as you select clients in the list.

Restricting a Client to Read-Only Access

To prevent a user on a client computer from modifying data on a SAN volume, you can mount the volume on the client with read-only access.



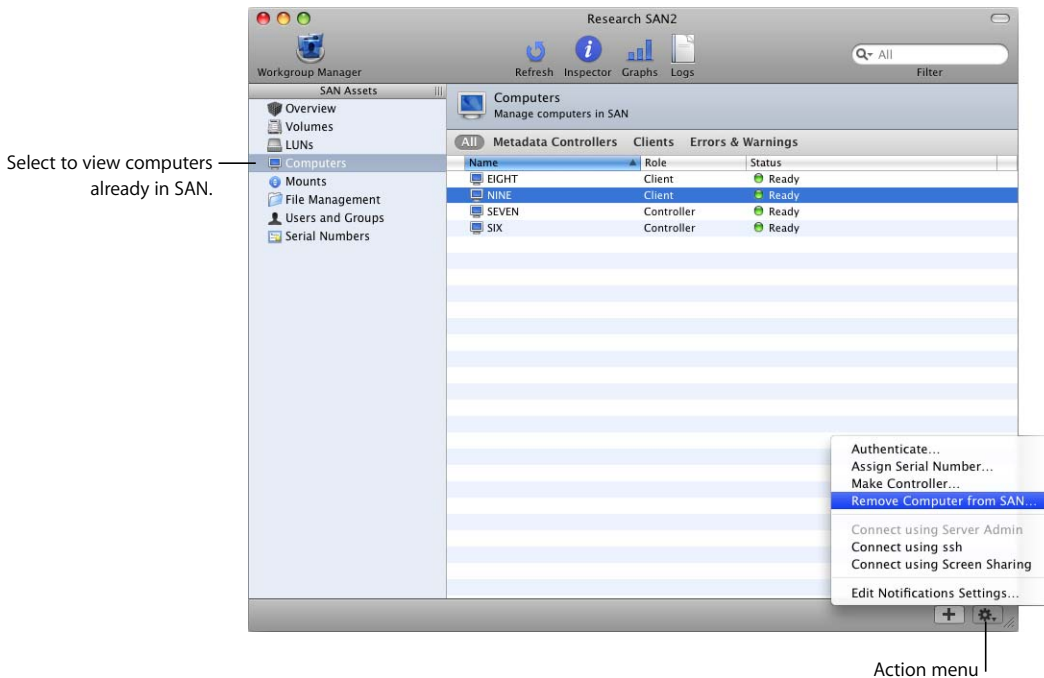
To mount a volume with read-only access:

- 1 In Xsan Admin, select Mounts in the SAN Assets list.
- 2 Choose the volume from the Volume pop-up menu.
- 3 Select the client.
- 4 Choose Mount Read Only from the Action (gear) pop-up menu.

If the volume is already mounted on the client, Xsan Admin unmounts it and mounts it again with read-only access.

Removing a Client from a SAN

You can remove a client computer from a SAN to prevent it from accessing SAN volumes.



To remove a client from a SAN:

- 1 In Xsan Admin, select Computers in the SAN Assets list, select the client, and choose "Remove computer from SAN" from the Action (gear) pop-up menu.
If SAN volumes are mounted on the client, Xsan Admin unmounts them automatically.
- 2 To remove the Xsan software from the computer, insert the *Xsan Install Disc* in the client computer and double-click Uninstall Xsan.pkg.
- 3 To prevent any possible connection to the SAN, physically disconnect the client computer from the SAN's Ethernet and Fibre Channel networks.

Mapping Windows User and Group IDs

You can use the Windows ID Mapping setting for a volume to specify how Windows clients map user and group information to Xsan-compatible user IDs (UIDs) and group IDs (GIDs), which they need in order to access Xsan volumes.

Note: To use ID mapping, Windows clients must be running StorNext 2.7. Xsan controllers and clients must be running Xsan 1.4.2 or later.

Windows clients can use either of these methods to provide UIDs and GIDs:

- **Generate IDs from GUID:** Windows clients dynamically generate UIDs and GIDs based on Globally Unique Identifier (GUID) information in an Active Directory domain. Choose this method for Macintosh computers on the SAN that are bound (connected) to Active Directory with the binding options set to automatically generate IDs.
- **Use IDs from LDAP (RFC 2307):** Windows clients get UID and GID values from the uidNumber and gidNumber attributes in Active Directory records. Choose this method for Macintosh computers on the SAN that are bound to Active Directory with the binding options set to map IDs to uidNumber and gidNumber.

The Active Directory binding options are set with Directory Utility in Mac OS X v10.5 Leopard or with Directory Access in Mac OS X v10.4 Tiger. These applications are in /Applications/Utilities/.

Important: To avoid ID conflicts, be sure all computers on the SAN use the same Active Directory domain and the same method of ID mapping.

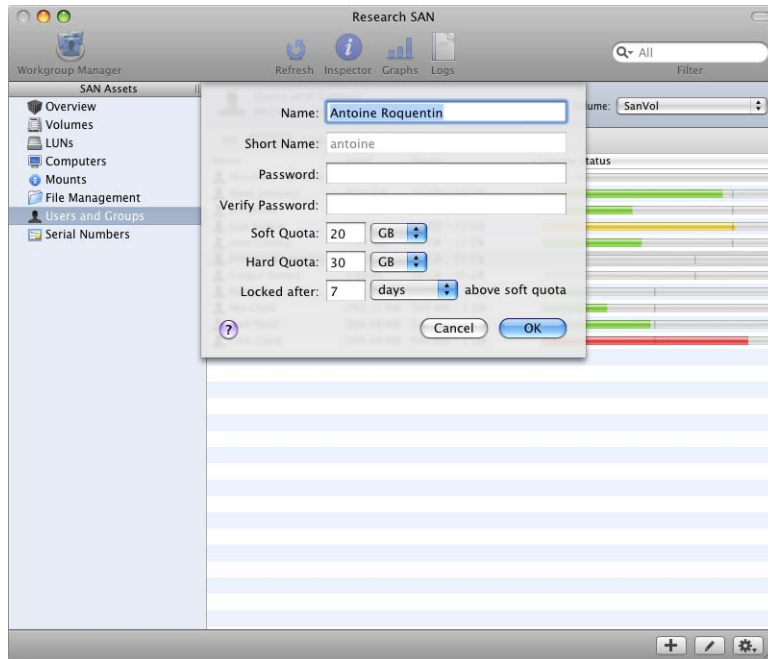
To select the Windows ID mapping method:

- 1 In Xsan Admin, select Volumes in the SAN Assets list, and choose Edit Volume Settings from the Action (gear) pop-up menu.
- 2 Choose a mapping method from the Windows ID Mapping pop-up menu.
If you choose “Use IDs from LDAP (RFC 2307)” you can change the ID numbers used when a directory record doesn’t include a uidNumber or gidNumber attribute.
- 3 Click OK.

Xsan Admin automatically unmounts the volume from all clients and controllers and stops the volume before changing the Windows ID mapping method, and then starts the volume and mounts it on each computer that had it mounted.

Setting SAN User and Group Quotas

You can use Xsan Admin to set quotas to manage the amount of storage available to a user or group.



To set a storage quota for a user or group:

- 1 In Xsan Admin, select Users and Groups in the SAN Assets list.
If you're not using Xsan Admin to manage users and groups, you'll see Quotas in the SAN Assets list instead of Users and Groups.
- 2 Choose a volume from the Volume pop-up menu.
- 3 Select a user or group in the list.
To select multiple users or groups, hold down the Shift or Command key as you select users or groups in the list.
To add a user or group, click the Users button or Groups button above the list and then click the Add (+) button.
- 4 Click the Edit button.
- 5 Enter a hard quota, soft quota, and grace period, and then click OK.

If your Xsan computers connect to another Mac OS X Server for user and group accounts, use Workgroup Manager or Server Preferences to create users and groups as needed. If existing users and groups aren't listed when you click the Add button, open Directory Utility (in /Applications/Utilities/) on your computer and make sure it's connected to the correct server for authentication. All computers in the SAN should use the same directory service.

From the Command Line

You can also set user quotas using the `cvadmin quotas set` command in Terminal. For information, see the `cvadmin` man page or "Viewing or Changing Volume and Storage Pool Settings (cvadmin)" on page 121.

About Xsan Quotas

Xsan enforces two disk space quotas for each user or group you choose to restrict: a soft quota and a hard quota. You can set these in combination to establish clear limits on the amount of storage a user or group can use, while still allowing temporary access to extra space for unexpected storage needs. You specify quotas individually for each volume on a SAN. A user who has no quotas specified can use all available space on a volume.

Soft Quota

The soft quota is the maximum space a user or group is expected to occupy on a regular basis. It is "soft" because it can be exceeded by an amount up to the hard quota for a grace period that you specify.

Hard Quota

The hard quota is an absolute limit on the space a user or group can occupy. Users are prevented from using more space than specified by their hard quotas.

Grace Period

A user or group can exceed the soft quota without penalty as long as each returns below the soft quota within the grace period you specify.

Soft Quotas Change to Hard Quotas

If a user or group exceeds the soft quota for a time longer than the grace period, the soft quota is changed to a hard quota. The user or group will not be able to save additional data on the volume until the user or group members delete enough old files to bring their usage below the soft quota.

Example

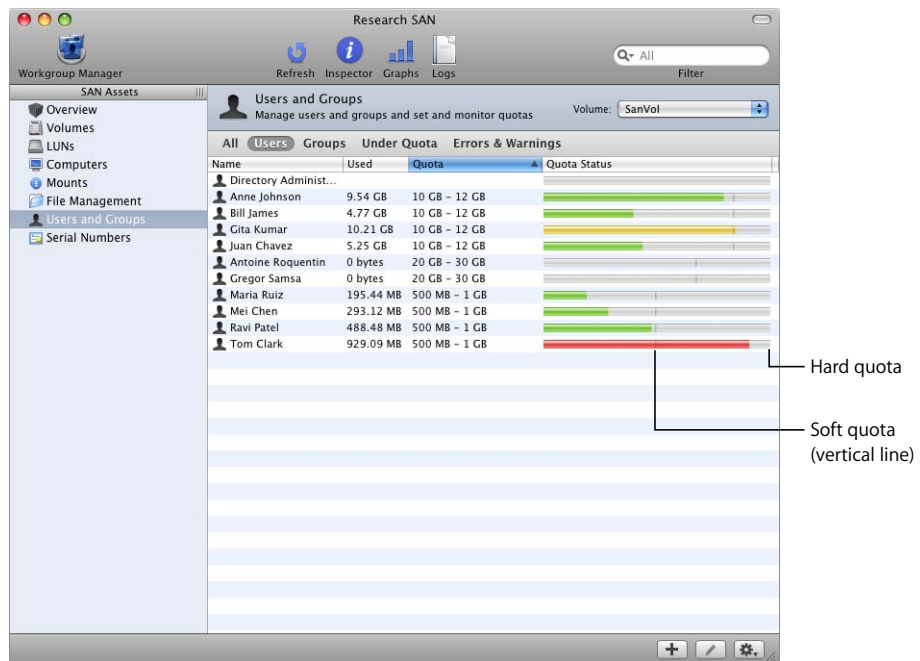
Suppose you assign Aldo a soft quota of 75 GB, a hard quota of 100 GB, and a grace period of 48 hours. Aldo's files can occupy up to 75 GB of space at any time, for as long as he needs them. If Aldo is surprised by additional or unusually large files, he can still copy them to the volume, up to a total of 100 GB. He then has 48 hours to remove files and return below the 75 GB soft limit. If he is still using more than 75 GB after 48 hours, Xsan resets his hard quota to 75 GB and he is forced to reduce his storage use. Aldo is unable to copy or save additional files to the volume until he deletes enough to return below the 75 GB quota.

Define SAN Users Consistently for Accurate Quotas

To be sure that Xsan user quota information is accurate, make sure user names and IDs are consistent on all computers on the SAN. For more information, see "Set Up SAN Users and Groups" on page 43.

Checking User Quota Status

You can use Xsan Admin to check file system quotas to see how much of their allotment users and groups are using.



To view quota status:

- In Xsan Admin, select “Users and Groups” or Quotas in the SAN Assets list. (You see Users and Groups only if you chose to have Xsan Admin manage your users and groups. Otherwise, you see Quotas instead.)

To be sure you’re seeing current information, click Refresh at the top of the window.

Xsan Admin displays the following information for each user or group:

- **Used:** The amount of space the user’s files are occupying.
- **Quota:** The soft and hard quotas. For example, “75 MB – 100 MB” indicates a soft quota of 75 MB and a hard quota of 100 MB.
- **Quota Status:** The status bar represents the full allocation, from zero on the left to the hard quota on the right. The small vertical line within the bar marks the soft quota. The colored portion of the bar shows how much space the user or group is currently using. Green indicates that the user or group is below the soft quota. Yellow indicates usage exceeding the soft quota but for a time within the grace period. Red indicates that the user has reached the hard quota, possibly because the soft quota was exceeded beyond the grace period and was changed to a hard quota.

You can set up Xsan to notify you by email or text message when a user or group exceeds a specific percentage of the quota. See “Setting Up Status Notifications” on page 105.


For more information about quotas and how to set them, see “Setting SAN User and Group Quotas” on page 85.

From the Command Line

You can also check user quotas using the `cvadmin quotas get` command in Terminal. For more information, see the `cvadmin` man page or “Viewing or Changing Volume and Storage Pool Settings (cvadmin)” on page 121.

Helping Clients Check Their Own Quotas

SAN users who work on client computers but don't have access to Xsan Admin can use the Xsan User Quotas application to check their own quotas.



To check your quota from a client computer:

- Open the Xsan User Quotas application (in /Applications/Server/).

Creating Local Home Folders for Network Accounts

Using a centralized directory simplifies the job of managing user accounts for client computers attached to the SAN. However, some applications, such as Final Cut Pro, work best when a user has a local home folder on the client he or she is using. User accounts that you manage with Xsan Admin are automatically set up with local home folders. If your SAN users have accounts on another directory system and they have network home folders, you can set up local home folders for them.

To create a local home folder for a network user account:

- 1 Open Workgroup Manager and authenticate to your SAN's Open Directory master.
- 2 Click the Users tab, select a user, and click Home.
- 3 If /Users appears in the list of home locations, select it and click Save.

If /Users is not in the list, click the Add (+) button. Then enter the following in the Full Path field (replacing *shortname* with the user's short name):

/Users/shortname

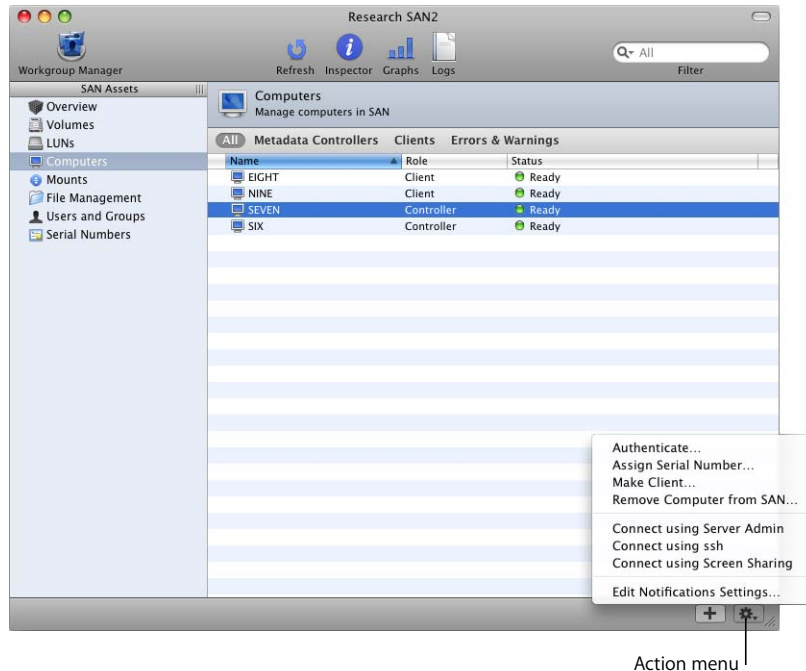
Leave all other fields blank, click OK, and then click Save.

The user's home folder is created on the client the first time the user logs in.

Accessing Client Computers Remotely

Xsan Admin can help you connect to an Xsan client computer so you can observe or control it over the network. Using Xsan Admin, you can:

- Start a screen sharing session so you can observe or control another computer.
- Open Terminal so you can log in using SSH and control another computer.
- Connect to and manage another server on the network using Server Admin.



Controlling a Client Using Screen Sharing

You can use the screen sharing feature of Mac OS X v10.5 Leopard to view and control the screen of a SAN client over the network. Xsan Admin can start a screen sharing session with the client.

To connect to a client using Leopard screen sharing:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the client you want to observe or control.
- 3 Choose “Connect Using Screen Sharing” from the Action (gear) pop-up menu.

If you have trouble sharing the screen of a remote computer, check the Sharing pane of System Preferences on the remote computer and make sure Remote Management service is turned on. If you have trouble starting a screen sharing session with an Xsan client, open Security preferences on the client, click Firewall, and make sure the option to block all incoming connections isn't selected.

Connecting to a Client Using SSH in Terminal

You can use the Secure Shell (SSH) tool from the command line to log in to a SAN client over the network. Xsan Admin can start an SSH session with the client or controller.

To connect to a client using SSH:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the client you want to connect to.
- 3 Choose "Connect using ssh" from the Action (gear) pop-up menu.

If you have trouble making an SSH connection, check the Sharing pane of System Preferences on the remote computer and make sure Remote Login service is turned on. If you have trouble making an SSH connection with an Xsan client, open Security preferences on the client, click Firewall, and make sure the option to block all incoming connections isn't selected.

Managing a Client Server Using Server Admin

You can remotely manage Mac OS X Server on a client by using the Server Admin application. Xsan Admin can open a Server Admin connection to a remote client server.

To connect to a controller or a client server using Server Admin:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the client you want to connect to.
- 3 Choose "Connect using Server Admin" from the Action (gear) pop-up menu.

To increase SAN security and redundancy, you can add, switch, and monitor Xsan metadata controllers.

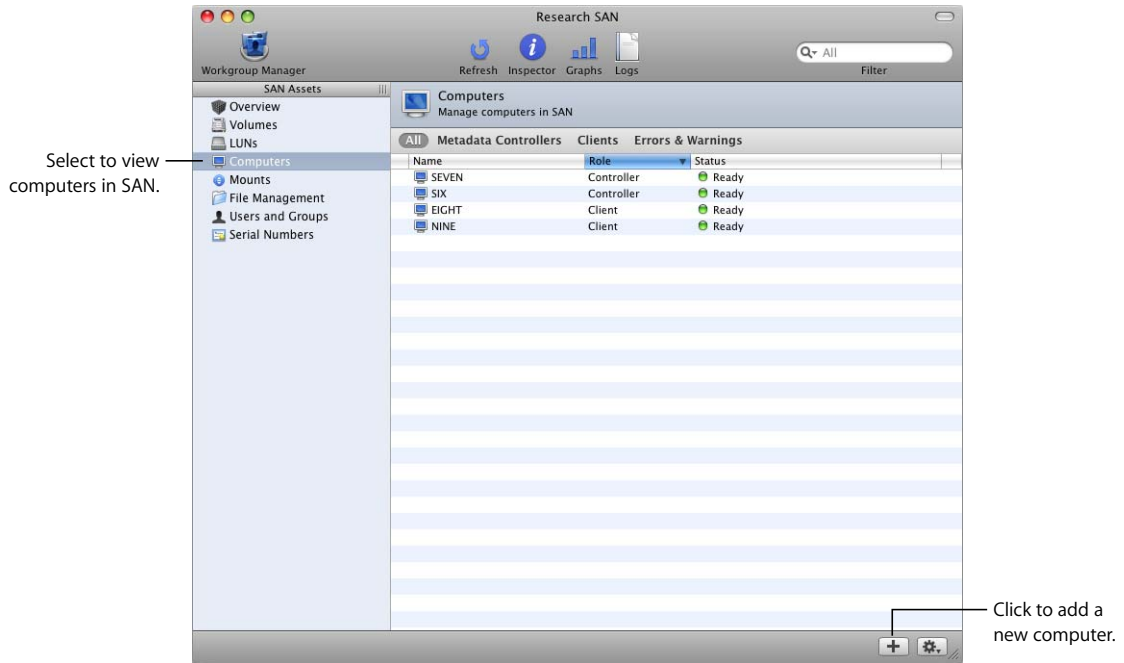
Every SAN volume you set up is managed by a metadata controller. To be sure that the volume is available to clients even if the primary metadata controller becomes unresponsive, you can set up standby controllers, one of which will assume control of the volume if the primary controller fails.

This chapter shows you how to add metadata controllers, set their failover priority, and force volume failover from the primary controller to a standby controller.

Adding a Metadata Controller

You can add one or more standby controllers to a SAN so that volumes are still available if the primary controller fails.

Any computer set to act as a metadata controller can also function as a client. So, if you don't want to dedicate a computer to act solely as a standby controller, you can use an existing client.



To add a metadata controller:

- 1 Connect the new controller computer to the SAN's Fibre Channel and Ethernet networks and install the Xsan software.
- 2 Open Xsan Admin, select Computers in the SAN Assets list, and then click the Add (+) button.
- 3 When the assistant opens, select the new controller computer in the computer list and click Continue.
If the computer doesn't appear in the list, click Add Remote Computer and add it.
- 4 On the Authenticate Clients pane, enter the administrator user name and password for the computer.
- 5 On the Choose Metadata Controllers pane, select the checkbox next to the computer in the list, and then click Continue.

Setting Controller Failover Priority

When the primary metadata controller for a volume fails, Xsan uses the failover priorities of the available standby controllers to decide which one to switch to.

To set a metadata controller's failover priority:

- 1 Open Xsan Admin, select Volumes in the SAN Assets list, and choose Edit Failover Priority from the Action (gear) pop-up menu.
- 2 Drag metadata controllers up or down in the list that appears. The closer a controller is to the top of the list, the more likely it is that it will host the volume.
- 3 Click OK.

Switching to a Standby Controller

You can use Xsan Admin to force an active metadata controller to turn over control of a volume to a standby controller.

To switch a volume to a standby metadata controller:

- 1 Open Xsan Admin and select Volumes in the SAN Assets list.
- 2 Select the volume in the list and choose “Force failover” from the Action (gear) pop-up menu.

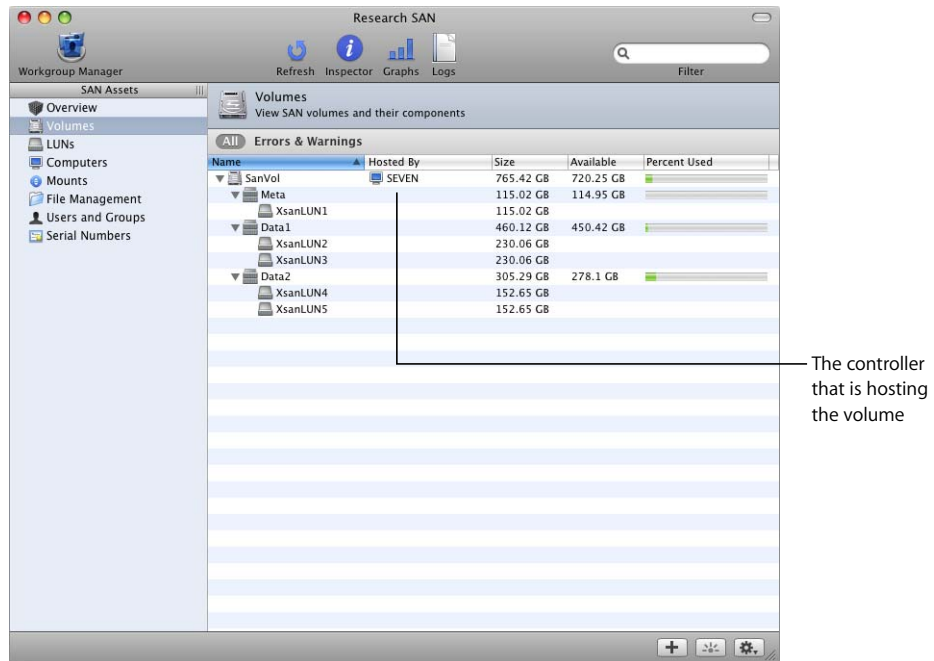
From the Command Line

You can also switch a volume to a standby metadata controller using the `cvadmin fail` command in Terminal.

For more information, see the `cvadmin` man page or “Viewing or Changing Volume and Storage Pool Settings (cvadmin)” on page 121.

Finding Out Which Controller Is Hosting a Volume

Control of a volume can move from one metadata controller to another as a result of controller failover. You can use Xsan Admin to find out which controller is currently hosting a volume.



To view a volume's controller:

- In Xsan Admin, select Volumes in the SAN Assets list and look in the Hosted By column.

From the Command Line

You can also find out which controller is hosting a volume by using the `cvadmin` command in Terminal. Open Terminal on the controller and type:

```
$ sudo cvadmin -e select
```

For more information, see the `cvadmin` man page or “Viewing or Changing Volume and Storage Pool Settings (cvadmin)” on page 121.

Listing the Volumes Hosted by a Controller

You can use Xsan Admin or the `cvadmin` command to find out which SAN volumes are being hosted by a particular metadata controller.

To list hosted volumes:

- In Xsan Admin, select Volumes in the SAN Assets list, and click the title of the Hosted By column to sort by controller.

From the Command Line

You can also find out which volumes are hosted by a controller using the `cvadmin select` command in Terminal. For more information, see the `cvadmin` man page or “Viewing or Changing Volume and Storage Pool Settings (cvadmin)” on page 121.

Changing a Controller’s IP Address

Follow these instructions if you need to change the IP address of an Xsan metadata controller.

WARNING: To avoid losing all data on the volumes hosted by the metadata controller, you must have a standby controller available.

To change a metadata controller’s IP address:

- 1 Make sure your standby controller is ready.
To check the status of the standby controller, open Xsan Admin and select Computers in the SAN Assets list.
- 2 In the Xsan Admin Computers pane, select the controller and choose “Remove Computer from SAN” from the Action (gear) pop-up menu.
- 3 Change the computer’s IP address in the Network pane of System Preferences.
- 4 Restart the computer.
- 5 With its new address, add the computer back to the SAN as a controller.
In Xsan Admin, select Computers in the SAN Assets list and click the Add (+) button.

If you want to switch control of a volume currently hosted by the standby controller back to the controller with the new IP address, select the volume in the Volumes pane of Xsan Admin and choose “Force failover” from the Action (gear) pop-up menu.

You can also force the switch from the command line by typing:

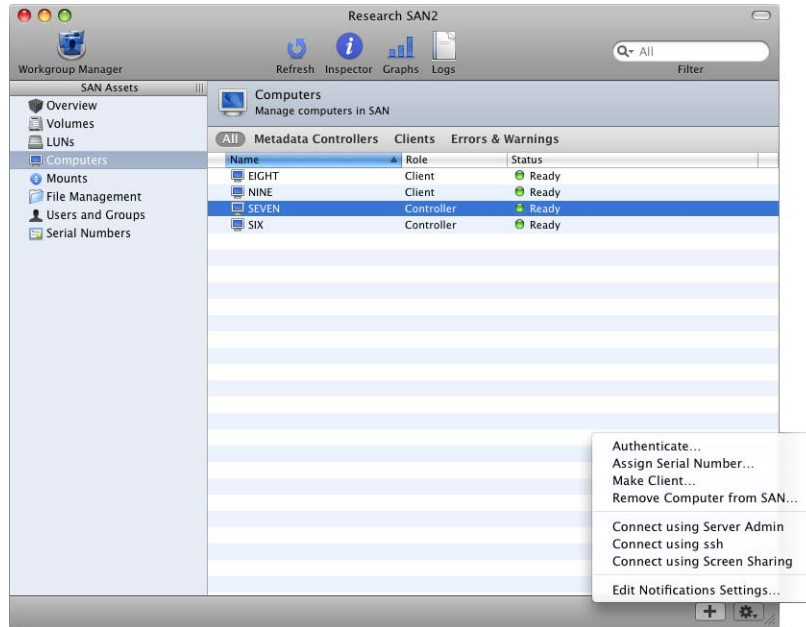
```
$ sudo cvadmin -e "fail volume"
```

where *volume* is the name of the Xsan volume.

Accessing Controller Computers Remotely

Xsan Admin can help you connect to an Xsan controller so you can observe or control it over the network. Using Xsan Admin, you can:

- Start a screen sharing session so you can observe or control another computer.
- Open Terminal so you can log in using SSH and control another computer.
- Connect to and manage another controller or client on the network using Server Admin.



Controlling a Controller Using Screen Sharing

You can use the screen sharing feature of Mac OS X v10.5 Leopard to view and control the screen of a SAN controller over the network.

To connect to a controller using Leopard screen sharing:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the controller you want to observe or control.
- 3 Choose “Connect using Screen Sharing” from the Action (gear) pop-up menu.

If you have trouble sharing the screen of a remote computer, check the Sharing pane of System Preferences on the remote computer and make sure Remote Management service is turned on. If the computer is running Mac OS X Server 10.5 Leopard, check that remote management is enabled in the General Settings pane of Server Admin.

Connecting to a Controller Using SSH in Terminal

You can use the Secure Shell (SSH) tool from the command line to log in to a SAN controller over the network.

To connect to a controller using SSH:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the controller you want to connect to.
- 3 Choose “Connect using ssh” from the Action (gear) pop-up menu.

If you have trouble making an SSH connection, check the Sharing pane of System Preferences on the remote computer and make sure Remote Login service is turned on. If the computer is running Mac OS X Server 10.5 Leopard, check that remote login is enabled in the General Settings pane of Server Admin.

Managing a Controller Using Server Admin

You can remotely manage Mac OS X Server on a controller by using the Server Admin application. Xsan Admin can open a Server Admin connection to a remote controller or client server running an Advanced configuration of Mac OS X Server.

To connect to a controller or a client server using Server Admin:

- 1 In Xsan Admin, select Computers in the SAN Assets list.
- 2 Select the controller you want to connect to.
- 3 Choose “Connect using Server Admin” from the Action (gear) pop-up menu.

Monitoring Controller Status

For information about checking or reporting the status of a controller, see these topics:

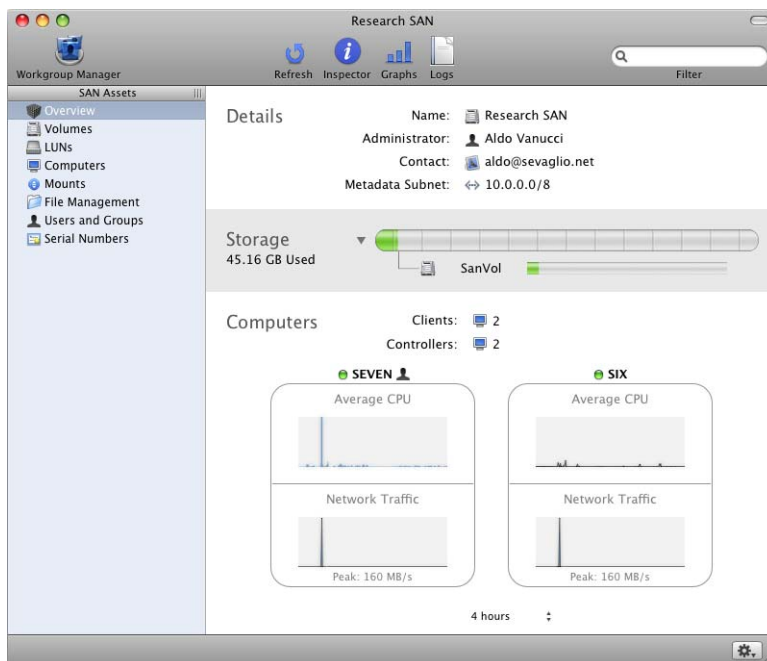
- “Graphing SAN Resource Usage” on page 104
- “Setting Up Status Notifications” on page 105
- “Viewing Xsan Logs” on page 106

You can use Xsan Admin and related command-line tools to check the condition of a SAN and its components.

This chapter shows you how to check the status of a SAN and its volumes and how to set up automatic notifications so you'll be alerted to changes in the SAN.

Checking SAN Status

You can use Xsan Admin to view status and configuration information for the SAN and its components.

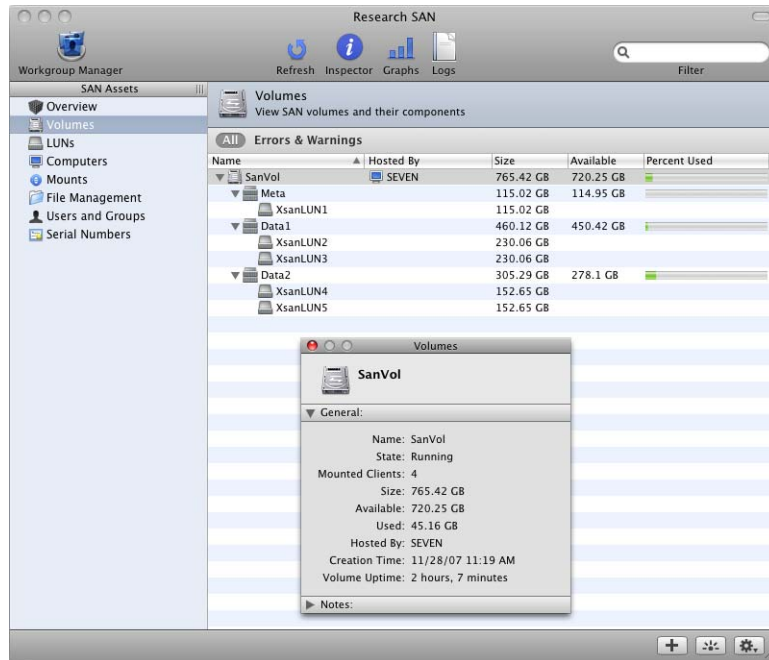


To view a component's status and configuration information:

- Open Xsan Admin, click the Inspector button at the top of the window, and then select the component in the SAN Assets list or the main pane of the Xsan Admin window.

Checking Volume Status

You can use Xsan Admin to view the status of a volume.

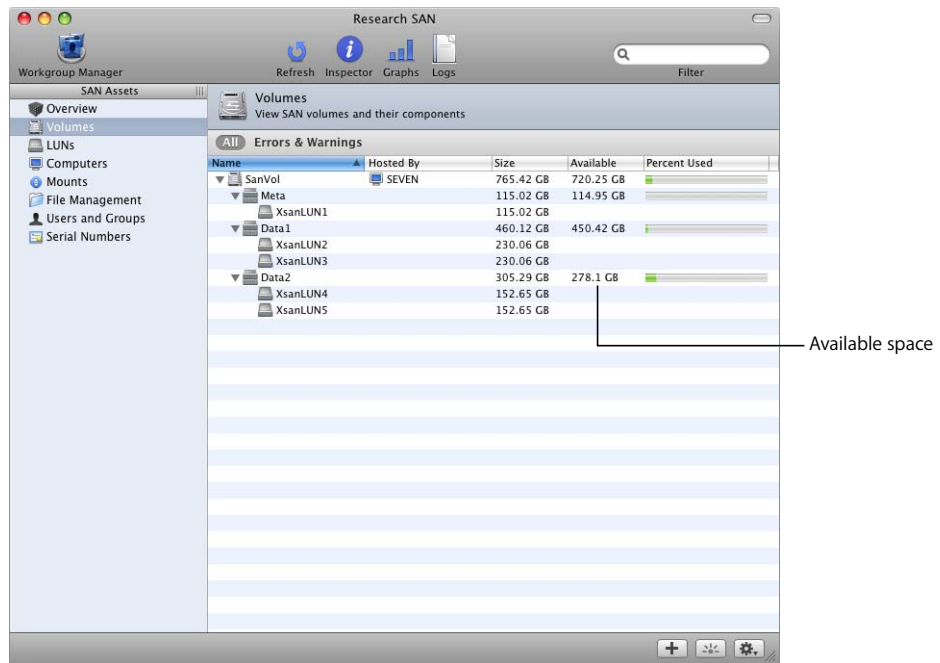


To get a quick look at the status of a volume:

- Open Xsan Admin, select Volumes in the SAN Assets list, select the volume, and click the Inspector button at the top of the window.

Checking Free Space on a Volume

There are several ways to see how much space is available on a SAN volume.



To check the free space on a volume:

- From a client or controller computer that has the volume mounted, select the volume in a Finder window and look at the size information at the bottom of the window (in Column or List view) or choose File > Get Info.
- From a computer that doesn't have the volume mounted or from a computer that doesn't belong to the SAN, open Xsan Admin, select the volume in the SAN Assets list, and click the Inspector button.

The reported size and free space for an Xsan volume doesn't include space in storage pools that contain only journal data and metadata. Only space on storage pools where users can store files is counted (that is, storage pools set to be used for "Any data" or "User data only"). For example, if you create a volume consisting of four 120 GB storage pools and configure one for journal and metadata only, Xsan Admin reports the size of the volume as 360 GB, not 480 GB.

From the Command Line

You can also check volume free space using the `cvadmin stat` command. For more information, see the `cvadmin` man page or "Viewing or Changing Volume and Storage Pool Settings (cvadmin)" on page 121.

Checking Free Space on a Storage Pool

In addition to checking free space on a volume, you can check free space on individual storage pools in the volume.

To check the free space on a storage pool:

- Open Xsan Admin, select Volumes in the SAN Assets list, select a storage pool, and then click the Inspector button at the top of the window.


If you don't see the storage pools for a volume, click the volume's disclosure triangle.

From the Command Line

You can also check storage pool free space using the `cvadmin show` command. For information, see the `cvadmin` man page or “Viewing or Changing Volume and Storage Pool Settings (cvadmin)” on page 121.

Graphing SAN Resource Usage

Xsan Admin can display graphs of up to a week of CPU, memory, Ethernet, and Fibre Channel utilization data for any computer on the SAN.



To view usage graphs:

- In Xsan Admin, click the Graphs button at the top of the window. Use the three pop-up menus in the Graphs window to choose a computer, a data type, and a time interval.

Memory and CPU resources used by the file system (fsm) process for a volume are listed under the name of the volume in the Graphs pop-up menu when you choose the volume's controller from the Computer pop-up menu.

From the Command Line

You can also check the file system process's current CPU and memory utilization by using the `top` command in Terminal to check the process named `fsm` on the volume's controller.

Setting Up Status Notifications

You can set up Xsan to send an email or dial a pager to notify you or other administrators when:

- A controller switches to its backup
- A Fibre Channel connection fails
- Free space on a volume falls below a specific percentage
- A user or group exceeds the designated soft quota
- A serial number is about to expire

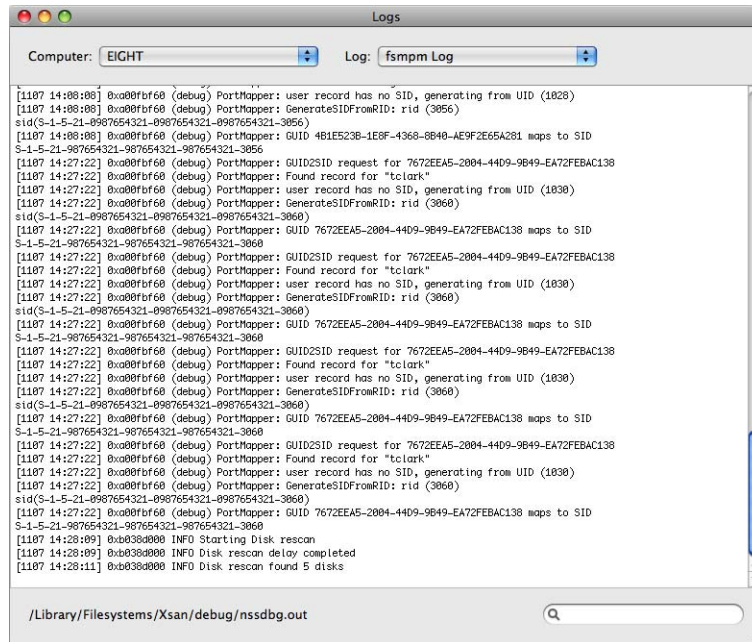
To send email notifications outside the local network, the controller needs access to an SMTP server.

To have Xsan send status notifications:

- 1 Open Xsan Admin and select Overview in the SAN Assets list.
- 2 Choose Edit Notifications Settings from the Action (gear) pop-up menu.
- 3 To add a new contact, click the Add (+) button and enter an email address.
- 4 If the address is for an account that will forward the notification as a text message, click the checkbox in the Text Msg column.
- 5 Choose the conditions that cause a notification to be sent (next to "Notify if").
- 6 Enter a sender name.
- 7 Enter the mail server address in the SMTP Server field.
- 8 To send a test message to all recipients, click Send Test Notification.
- 9 Adjust settings as necessary, and then click OK.

Viewing Xsan Logs

You can use Xsan Admin to view the informational and diagnostic messages that Xsan writes to a computer's system and log.



To view the SAN logs:

- In Xsan Admin, click the Logs button at the top of the window. In the Logs window that appears, use the Computer and Log pop-up menus to choose the log you want to view.

To display just entries containing specific names, times, or other text, type in the Search field in the lower-right corner of the window.

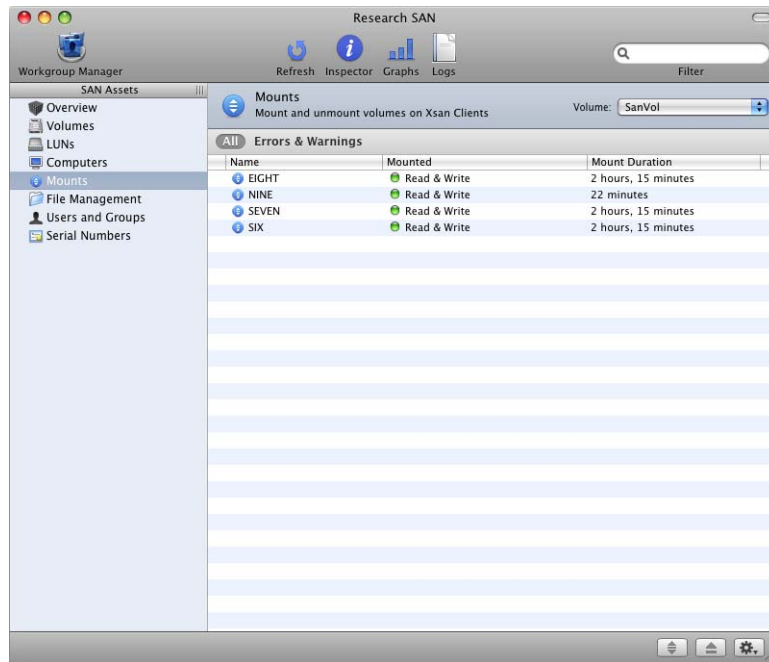
From the Command Line

To see the log for a particular volume from the command line, look at the log file:

/Library/Filesystems/Xsan/data/volume/log/cvlog

Checking Volume Clients

You can use Xsan Admin to see a summary of which clients are using a volume.



To see how many clients have a volume mounted:

- Open Xsan Admin and select Volumes in the SAN Assets list. Select the volume in the list and click the Inspector button at the top of the window.

To see which clients are using a volume:

- Open Xsan Admin, select Mounts in the SAN Assets list, and choose the volume from the Volume pop-up menu.

From the Command Line

You can also use the `cvadmin who` command in Terminal to see a list of volume clients. For more information, see the `cvadmin` man page or “Viewing or Changing Volume and Storage Pool Settings (cvadmin)” on page 121.

Checking for Fibre Channel Connection Failures

Fibre Channel connection failures or errors are recorded in the system log.

To view the system log:

- In Xsan Admin, click the Logs button at the top of the window. In the Logs window that appears, choose the computer in the Computer pop-up menu and choose System Log from the Log pop-up menu.

This chapter provides solutions to common problems you might encounter while working with a SAN.

Look here for solutions to common problems you might encounter while setting up, managing, or using an Xsan SAN.

If You're Unable to Connect to a Computer Using Xsan Admin

If there is a firewall between the administrator computer and the SAN computer, make sure TCP port 311 is open.

If You're Unable to Install the Xsan Software

If the installer says you can't install the Xsan 2 software on a particular computer, make sure the computer has Mac OS X v10.5 Leopard or Mac OS X Server v10.5 Leopard installed.

If Some Computers Aren't Listed in Xsan Admin

If a computer you want to add to the SAN as a metadata controller or client is not listed in Xsan Admin, make sure:

- You have installed the Xsan software on the computer.
- The computer is powered on.
- The computer is not sleeping and is not set to sleep (in the Energy Saver pane of System Preferences).
- The computer is on the same TCP/IP subnets as the other SAN components. (If you are using both a private and a public Ethernet network, all SAN components must be connected to both.)

If You're Unable to Mount a Volume on a Client

- Try restarting the client computer, and then try again.
- Check that all Fibre Channel cables are plugged in.
- Make sure there are no other volumes mounted on the client that have the same name as the Xsan volume.

If RAID LUNs Aren't Accessible over Fibre Channel

- Try restarting the computer that doesn't see the LUNs.
- Check the configuration of the Fibre Channel switch to be sure the SAN components are in the same Fibre Channel zone.

If Files and Folders Created by Mac OS 9 Computers Show the Wrong Creation Date

A computer running Mac OS 9 can store files or folders on an Xsan volume that is shared using the AFP service of Mac OS X Server. However, the creation date reported for those items by the Finder's Get Info command on the Mac OS 9 computer is always Feb 4, 2040. A computer running Mac OS X reports a blank creation date for the same files and folders. In both cases, the correct creation date is displayed in the Modified field.

If You Have Problems Using Command-Line Tools

- If you get the response "Cannot list FSS - reason -Bad file descriptor" when you run the `cvadmin` tool, make sure you have root user privileges when you use the tool. Either log in as the root user or use the `sudo` command to run the tool. For example:

```
$ sudo cvadmin
```

If a LUN Doesn't Have as Much Space as Expected

To make striping across LUNs possible, Xsan automatically adjusts LUN sizes to make all LUNs in a storage pool the same size as the smallest LUN in the pool. Xsan doesn't use the extra space on larger LUNs when you mix LUNs of different sizes in a storage pool.

If You're Unable to Rename an Xsan Volume in the Finder

Xsan doesn't let a mounted Xsan volume be renamed using the Finder. To rename an Xsan volume, reinitialize the volume with Xsan Admin. For instructions, see "Renaming a Volume" on page 65.

If You're Unable to Add a Storage Pool

There are some reserved names you can't use to name a storage pool. If you enter one of these names, the OK button in the storage pool sheet is disabled.

Reserved storage pool names		
Affinity	InodeExpandInc	Rtios
AllocationStrategy	InodeExpandMax	RtiosReserve
AttrTokenSize	InodeExpandMin	Rtmb
Brls	IoHangLimitSecs	RtmbReserve
BrlTime	Journal	RtTokenTimeout
BufferCacheSize	JournalLcBufNum	Sectors
BufferPoolSize	JournalLcBufSize	SectorSize
BWMFields	JournalSize	Static
DataMigration	Log	StaticNodes
DataMigrationThreadPoolSize	MaxConnections	Status
Debug	MaxLogs	Sticky
DeviceName	MaxLogSize	StripeBreadth
DirCacheSize	MaxMBPerClientReserve	StripeClusters
DirFDCacheSize	Mbufs	StripeGroup
DirWarp	MbufSize	ThreadPoolSize
Disabled	MetaData	Type
Disk	MirrorGroup	UnixFabricationOnWindows
DiskType	MirrorReadMethod	UnixNobodyUidOnWindows
Enabled	MultiPathMethod	UnixNobodyGidOnWindows
Exclusive	No	UnixFileCreationMode ...
ForcePerfectFit	Node	UnixDirectoryCreation ...
ForceStripeAlignment	OpHangLimitSecs	WindowSecurity
FSBlockSize	Quotas	Write
GlobalSuperUser	Read	Yes
InodeCacheSize	Regular	
InodeDeleteMax	Rotate	

If Fibre Channel Performance Is Poorer Than Expected

Mismatched optical transceivers (GBICs) can cause Fibre Channel communication errors and degrade SAN performance. To ensure good performance, use identical transceivers (same manufacturer and model number) on both ends of your Fibre Channel cables.

If a Client is Unable to Use a Volume After a Fibre Channel Interruption

If a client loses its Fibre Channel connection to the SAN (because a cable is unplugged, for example), the client might not recognize LUNs in an Xsan volume after the connection is restored. If this happens, restart the client to remount the volume.

If problems persist, restart all SAN devices. Restart RAID systems first, then continue with SAN controllers, and finally, restart all clients.

To check whether a computer is seeing Xsan volume LUNs:

- Open Disk Utility on the computer and look for the LUNs in the list of disks and volumes.

From the Command Line

You can also check for accessible LUNs using the `cvlabel -l` command or the `diskutil list` command in Terminal.

If You're Unable to Add LUNs to a Storage Pool

You can't add a new LUN to an existing storage pool unless the LUN is at least as large as the common LUN size for the pool. (The common LUN size is the size of the smallest LUN you added when you created the pool.) You can add a larger LUN, but space beyond the common LUN size isn't used.

You can only expand storage pools that can be used for user data. You can't add a LUN to an existing storage pool if the storage pool can be used only for journaling and metadata. If you want to add journaling and metadata storage, add another storage pool that can be used for journaling and metadata.

To check the common LUN size for a storage pool:

- In Xsan Admin, select Volumes in the SAN Assets list, and then click disclosure triangles in the list of volumes to show the LUNs in the storage pool of interest. Compare the listed LUN sizes.

To check the size of the LUN you want to add:

- In Xsan Admin, select LUNs in the SAN Assets list, and then click Unused LUNs. Check the listed size for the LUN of interest.

To check the data types a storage pool is used for:

- In Xsan Admin, select Volumes in the SAN Assets list, and then click disclosure triangles in the list of volumes to show the storage pool of interest. Double-click the storage pool in the list, and look next to "Used For" in the Inspector window.

If the Capacity of a Larger LUN is Listed as 2 Terabytes

If a LUN that doesn't yet belong to a storage pool is listed in Xsan Admin with a capacity of 2 TB, even though you know it is larger (which can happen if you used the LUN with an earlier version of Xsan), try relabeling the LUN.

To relabel a LUN:

- 1 In Xsan Admin, select LUNs in the SAN Assets list.
- 2 Click LUNs, select the LUN in the list, and choose Remove LUN Label from the Action (gear) pop-up menu.
- 3 With the LUN still selected, choose Change LUN label from the Action pop-up menu, and enter a label.

If File Copying Doesn't Finish

If the Ethernet connection to a metadata controller is lost, Finder file-copy operations in progress on clients might not finish, even though the volume successfully fails over to a standby controller.

To allow the copy operation to finish:

- Reconnect the disconnected controller to the SAN's Ethernet network.

Combining Xsan Controllers and StorNext Clients

This appendix shows how to connect Windows, Solaris, UNIX, AIX, IRIX, or Linux clients to an Xsan SAN.

Xsan is fully compatible with Quantum’s StorNext File System, so you can set up Xserve and RAID systems to act as SAN controllers and storage for Windows, Sun Solaris, UNIX, IBM AIX, SGI IRIX, or Linux clients that are running StorNext FX software. For information about adding Macintosh clients to an existing StorNext SAN, see the StorNext documentation.

Terminology

Note these differences in terminology between StorNext and Xsan:

StorNext term	Equivalent Xsan term
file system	volume
file system server (FSS)	controller (or metadata controller)
stripe group	storage pool

Compatible Software Versions

This table shows which versions of Xsan and StorNext controllers and clients can be used on the same SAN.

Controller	Client	Compatible
Xsan 2	StorNext FX 1.4 or 2.0	Yes
	StorNext FX 1.3	No
	StorNext FS/FX 2.4–3.1	No
StorNext FS 3.1	Xsan 2	Yes
	Xsan 1.4.2	Yes
StorNext FS 3.0	Xsan 2	No
	Xsan 1.4.2	Yes
StorNext FS 2.8	Xsan 1.4.1–1.4.2 and Xsan 2	No
	Xsan 1.4	Yes
StorNext FS 2.4–2.7	Xsan 1.4–1.4.2 and Xsan 2	No

Licensing

You license a Macintosh Xsan client or controller using either the single-copy serial number printed on the *Xsan Install Disc* sleeve that comes in the Xsan package, or serial numbers you purchase separately. Licenses for StorNext are purchased from Quantum when you buy the StorNext software. Xsan clients do not use or count against StorNext File System client licenses.

Using Xsan Controllers with StorNext Clients

You can use Quantum's StorNext software to access an Xsan SAN from a Windows, UNIX, Sun Solaris, IBM AIX, SGI IRIX, or Linux computer.

- 1 Connect the non-Macintosh client to the SAN's Fibre Channel and Ethernet networks.
- 2 Install the StorNext File System software on the non-Macintosh client following the instructions that Quantum provides in the StorNext package.
- 3 Duplicate the Macintosh Xsan controller's shared secret file on the non-Macintosh client.

The shared secret file is named `.auth_secret`. On a Macintosh Xsan controller, it is stored in the folder `/Library/Filesystems/Xsan/config/`.

Copy the file (using the same name) to the non-Macintosh client. On SGI IRIX, Sun Solaris, IBM AIX, and Linux StorNext clients, put the file in `/usr/cvfs/config/`. On Windows clients, put the file in `%cvfsroot%\config\`, where `%cvfsroot%` is the folder where you installed StorNext.

Important: This file contains sensitive information. Secure the file for read/write access by the root user or Windows administrator only.

- 4 Place a StorNext license file for your non-Macintosh clients on the Macintosh Xsan controller.

On the Xsan controller, put the file (named `license.dat`) in the folder:

`/Library/Filesystems/Xsan/config/`

Contact Quantum to obtain a license file for your non-Macintosh clients.

This appendix describes Xsan shell commands and configuration files you can use to work with a SAN from the command line.

You can use the shell commands and configuration files described here to access, set up, and manage Xsan SANs, LUNs, storage pools, and volumes from the command line. The Terminal application is the Mac OS X gateway to the BSD command-line interface (UNIX shell command prompt). Each window in Terminal contains a complete command-line execution context, called a shell, that is separate from all other execution contexts. Although you can use any shell of your choice, the examples in this book assume that you are using `bash`, the standard Mac OS X shell.

Using the Shell Commands

The Xsan command-line utilities are located in `/Library/Filesystems/Xsan/bin/`, which is part of the default shell search path.

Many commands used to manage Xsan must be executed by the root user (also known as the superuser). If you get a message such as “permission denied,” the command probably requires root user privileges.

To execute a single command with root user privileges, begin the command with `sudo` (short for superuser do). For example:

```
$ sudo cvfsck -n MyVolume
```

If you haven’t used `sudo` recently, you’re prompted for the password for your administrator account.

Sending Commands to Remote Computers

To use commands on a remote computer, first use SSH to log in to the other computer:

```
$ ssh user@computer
```

where *user* is a user account on the remote computer and *computer* is its IP address or DNS name.

Viewing the Man Pages

Detailed documentation for Xsan command-line utilities is available in UNIX-style man pages. A command's man page includes information about the command, its options, parameters, and proper use. The man pages for Xsan commands are located in `/Library/Filesystems/Xsan/man/`.

To view a man page, type:

```
$ man command
```

where *command* is the command you want information about.

Notation Conventions

These conventions are used throughout the command descriptions:

Notation	Indicates
fixed-width font	A command or other text typed in a Terminal window
\$	A shell prompt
[text_in_brackets]	An optional parameter
(one other)	Alternative parameters (type one or the other)
<i>italicized</i>	A parameter you must replace with a value
[...]	A parameter that may be repeated
<angle_brackets>	A displayed value that depends on your SAN configuration

Installing Xsan from the Command Line

- 1 Log in to a computer that does have a keyboard and monitor, and then insert the *Xsan Install Disc*.
- 2 Open the Terminal application (in `/Applications/Utilities/`).
- 3 In Terminal, copy the Xsan installer package to the remote computer:

```
$ scp -r /Volumes/Xsan\ Install\ Disc/Install\ Xsan.mpkg user@remotehost:/tmp/
```

where *user* is an administrator user on the remote computer and *remotehost* is the IP address or DNS name of the computer you want to install on.

If you want to install just the Xsan Admin application, enter:

```
$ scp -r /Volumes/Xsan\ Install\ Disc/Other\ Installs/XsanAdmin.mpkg/  
user@remotehost:/tmp/
```

- 4 Log in to the remote computer:

```
$ ssh user@remotehost
```

where *user* and *remotehost* are the same as in the previous step.

5 Run the installer on the headless computer:

- If you want to install the Xsan file system and Xsan Admin application, enter:

```
$ sudo installer -pkg /tmp/Install\ Xsan.mpkg -target /
```

- If you want to install just the Xsan file system, enter:

```
$ sudo installer -pkg /tmp/Install\ Xsan.mpkg/Contents/Installers/  
XsanFileSystem.mpkg/ -target /
```

- If you want to install just the Xsan Admin application, enter:

```
$ sudo installer -pkg /tmp/Install\ Xsan.mpkg/ -target /
```

If you want to watch the progress of installation, add the `-verbose` parameter:

```
$ sudo installer -verbose -pkg /tmp/Install\ Xsan.mpkg -target /
```

Xsan Commands

Xsan includes these command-line tools:

Tool	Description
<code>cvadmin</code>	View or change volume and storage pool settings (page 121)
<code>cvaffinity</code>	Manipulate affinity tags manually (page 124)
<code>cvcp</code>	Copy files or folders (page 124)
<code>cvfsck</code>	Check or repair a volume (page 125)
<code>cvlabel</code>	View, label, and initialize LUNs (page 126)
<code>cvmkdir</code>	Create a folder and assign an affinity (page 127)
<code>cvmkfile</code>	Create and preallocate a file (page 127)
<code>cvmkfs</code>	Initialize a volume (page 128)
<code>cvupdatefs</code>	Apply volume setup changes (page 128)
<code>snfsdefrag</code>	Defragment a volume (page 128)
<code>xsanctl</code>	Mount and unmount Xsan volumes (page 130)

Viewing or Changing Volume and Storage Pool Settings (cvadmin)

Use the `cvadmin` tool to perform status and setup tasks related to Xsan volumes.

For help, see the `cvadmin` man page or type:

```
$ sudo cvadmin -e help
```

To enter interactive mode:

```
$ sudo cvadmin
```

To execute commands from a file:

```
$ sudo cvadmin [-H host] [-F volume] -f cmdfile
```

To execute a single command and return to the shell prompt:

```
$ sudo cvadmin [-H host] [-F volume] -e ["command [cmdparam..."]]
```

Parameter	Description
<code>-H <i>host</i></code>	The metadata controller that is hosting the volume. If not provided, the local computer is assumed. <i>host</i> – the IP address or DNS name of a metadata controller other than the one on which you are logged in.
<code>-F <i>volume</i></code>	The volume to be the active ("selected") volume in <code>cvadmin</code> . <i>volume</i> – the name of an Xsan volume
<code>-f <i>cmdfile</i></code>	Read commands from the specified file. <i>cmdfile</i> – the name of a text file containing <code>cvadmin</code> commands
<code>-e <i>command</i></code>	Execute the specified command and return to the shell prompt. Otherwise, <code>cvadmin</code> continues to run in interactive mode with the prompt <code>Xsanadmin></code> . If you include parameters (<i>cmdparam</i>) with the command, enclose the command and its parameters in a pair of quotes. Available commands are listed under "cvadmin Commands," below.
<i>cmdparam</i>	Values required by the command.

Commands available in the `cvadmin` tool are listed in the following table.

cvadmin Commands

cvadmin command	Description
<code>activate</code> <code>[<i>volume</i> <i>index</i>]</code>	Choose the "active" volume that you want to work with interactively. <i>volume</i> – the name of the volume <i>index</i> – the numeric ID of the volume (to see a list of these, use the <code>cvadmin select</code> command without any parameters)
<code>disks [refresh]</code>	List LUNs.
<code>down <i>pool</i></code>	Disallow all access to a storage pool. <i>pool</i> – the name of a storage pool in the currently active volume
<code>fail (<i>volume</i> <i>index</i>)</code>	Cause a volume to fail over to a standby controller. <i>volume</i> – the name of the volume <i>index</i> – the numeric ID of the volume (to see a list of these, use the <code>cvadmin select</code> command without any parameters)
<code>filelocks [yes no]</code>	Enable or disable file and record locks. Use the command without any parameter to see the current setting for locks.
<code>multipath <i>pool</i></code> <code>(rotate static)</code>	Specify how Xsan uses multiple paths to a storage pool. <i>pool</i> – the name of a storage pool in the currently active volume
<code>paths</code>	List available LUNs.
<code>quit</code>	Exit from <code>cvadmin</code> .

cvadmin command	Description
quotas [yes no]	Enable or disable quotas for the active (selected) volume. Use the command without any parameters to see the current setting for quotas.
quotas get (user group) <i>name</i>	Display current quota information for a user or group. <i>name</i> – the name of the user or group
quotas set (user group) <i>name hard soft grace</i>	Set quotas for user or group name. <i>name</i> – the name of the user or group <i>hard</i> – hard quota (bytes) <i>soft</i> – soft quota (bytes) <i>grace</i> – grace period (minutes)
quotacheck	Recalculate quota information for the active volume.
repquota	Generate the following quota report files in /Library/Filesystems/Xsan/data/ <i>volume</i> : quota_report.txt – text file quota_report.csv – comma-delimited file quota_regen.in – cvadmin commands that will set up identical quotas on another metadata controller. You can use cvadmin -f to execute the commands.
repof	Create a report of open files on the active volume in the file /Library/Filesystems/Xsan/data/ <i>volume</i> /open_file_report.txt.
select [<i>volume</i>]	Choose the “active” volume that you want to work with. The name of the currently active volume appears preceding the command prompt in interactive mode, for example: Xsanadmin (Vol1) > To see a list of running volumes, leave off the volume parameter. <i>volume</i> – the name of an Xsan volume
show [<i>pool</i>] [long]	List storage pool information for the active volume. <i>pool</i> – the name of a storage pool in the currently active volume
start <i>volume</i> [on] [<i>controller</i>]	Start a volume based on the information in its configuration file (/Library/Filesystems/Xsan/config/ <i>volume</i> .cfg). <i>volume</i> – the name of an Xsan volume <i>controller</i> – The address of the metadata controller to start the volume’s FSM process on
stat	Display information about the active volume.
stop <i>volume</i>	Stop a volume and its FSM process on all metadata controllers.
up <i>pool</i>	Allow access to the specified storage pool. <i>pool</i> – the name of a storage pool in the currently active volume
who	Display client information for the active volume.

Manipulating Affinity Tags (cvaaffinity)

Use the `cvaaffinity` command to assign an affinity tag to a folder or a file, or to list the affinity tag currently assigned to a folder or a file. Assigning an affinity tag to a folder or file causes it to be stored on a storage pool that has the same affinity tag. You can see the affinity tags for available storage pools by using the `show long` command of the `cvadmin` tool.

To set an affinity tag for a folder or file:

```
$ cvaaffinity -s affinity target
```

To list the affinity tag currently assigned to a folder or file:

```
$ cvaaffinity -l target
```

To delete the affinity tag from a folder or file:

```
$ cvaaffinity -d target
```

Parameter	Description
<i>affinity</i>	The affinity tag that's assigned to the storage pools where you want the target folder or file to be stored
<i>target</i>	The path to and name of the folder or file.

Copying Files or Folders (cvcp)

Use the `cvcp` command to copy files or folders to or from an Xsan volume.

```
$ cvcp [options] source destination
```

Parameter	Description
<i>options</i>	See "cvcp Command Options," below.
<i>source</i>	The file or folder (directory) to be copied.
<i>destination</i>	Where the copy is created.

cvcp Command Options

Option	Description
-A	Turn off preallocation.
-b <i>buffers</i>	Set the number of I/O buffers to use. <i>buffers</i> – the number of buffers to use for the copy
-k <i>size</i>	Set the copy buffer size. <i>size</i> – the buffer size (bytes)
-l	Copy the targets of symbolic links, not the links.
-n	Do not apply command to subfolders.
-p <i>prefix</i>	Only copy files with names that start with the specified prefix. <i>prefix</i> – characters to match with the beginning of the file name
-s	Allocate on storage pool block boundaries.
-t	Specify the number of copy threads.

Option	Description
-v	Report all information about file copied.
-x	Retain original file permissions in the copy.
-y	Retain ownership and group information in the copy. This works only if the root user is performing the copy.
-z	Retain original modification times in the copy.

Examples

Copy the file `friday` to `/datasets/data1/july/`:

```
$ cvcp friday /datasets/data1/july
```

Copy the folder `/data1/` and all subfolders to `/datasets/data1/`, retaining all permissions and ownerships and displaying files as they are copied:

```
$ cvcp -vxy data1 /datasets/data1
```

Perform a similar copy as above, but only copy files with names that begin “jul”:

```
$ cvcp -vxy -p jul data1 /datasets/data1/july
```

Checking or Repairing a Volume (cvfsck)

Use the `cvfsck` command to check or repair an Xsan volume.

```
$ sudo cvfsck [options] volume
```

Parameter	Description
<i>options</i>	See “cvfsck Command Options,” below.
<i>volume</i>	The name of the volume to check or repair.

cvfsck Command Options

Option	Description
-d	Display extra debugging information.
-e	Display file extents statistics.
-f	Report fragmentation.
-g	Print journal recovery log.
-j	Perform journal recovery.
-J	Display raw journal data.
-K	Reset journal. WARNING: Resetting the journal might introduce metadata inconsistencies. Don't use unless absolutely necessary.
-l	Record problems in the system log.
-n	Check volume in read-only mode.
-r	Relocate files before changing volume configuration.
-v	Display all available information.

Option	Description
-w	Modify the file system as needed to repair problems.
-x	Report statistics in comma-separated form for use in a spreadsheet.

Labeling, Listing, and Unlabeling LUNs (cvlabel)

Use the `cvlabel` command to initialize LUNs so they can be added to storage pools. For details, see the `cvlabel` man page.

To list available LUNs:

```
$ sudo cvlabel -l [-s] [-v]
```

To list current LUN and label information you can paste into a label file:

```
$ sudo cvlabel -c
```

To label a LUN:

```
$ sudo cvlabel [-v] [-f] [labelfile]
```

To remove the existing label from a LUN:

```
$ sudo cvlabel -u lun
```

Parameter	Description
-l	List available LUNs.
-s	Display device serial numbers.
-v	Show progress display.
-c	Create a label template file.
-f	Relabels LUNs that are already labeled.
<i>labelfile</i>	An optional file containing information for each label. You can use the -c option to create this file, or use this file as a template: /Library/Filesystems/Xsan/examples/cvlabels.example
<i>lun</i>	The LUN identified by disk name, for example /dev/disk4.
-u	Unlabel the specified LUN.

Creating a Folder and Assigning an Affinity (cvmkdir)

Use the `cvmkdir` command to create a folder (directory) and assign it an affinity tag so that its contents are stored on storage pools with the same affinity tag.

```
$ cvmkdir -k affinity folder
```

Parameter	Description
<i>-k affinity</i>	Specifies the affinity tag to be associated with the folder. <i>affinity</i> – the affinity tag that’s assigned to the storage pools where you want the folder’s contents to be stored You can use the <code>show long</code> command of the <code>cvadmin</code> tool to see a storage pool’s affinity tag. You can use <code>-k ""</code> to remove the folder’s affinity tag.
<i>folder</i>	The path to and name of the folder.

Creating and Preallocating a File (cvmkfile)

Use the `cvmkfile` command to allocate space for a file on an Xsan volume.

```
$ cvmkfile [-k affinity] [-p] [-s] [-w] [-z] size(k|m|g) filename
```

Parameter	Description
<i>-k affinity</i>	Allocate space for the file on one of the storage pools with the specified affinity tag. <i>affinity</i> – the affinity tag that’s assigned to the storage pools where you want the folder’s contents to be stored You can use the <code>show long</code> command of the <code>cvadmin</code> tool to see a storage pool’s affinity tag.
<i>-p</i>	Force future extensions of the file to be aligned on block boundaries.
<i>-s</i>	Force the file allocation to align with block boundaries.
<i>-w</i>	Set file size as indicated by <i>size</i> .
<i>-z</i>	Set the contents of the file to zeros.
<i>size(k m g)</i>	A number specifying the amount of space to allocate to the file. <i>size</i> – a number k – kilobytes m – megabytes g – gigabytes
<i>filename</i>	The path to and name of the file to allocate.

Example

Allocate 2 GB of space for the file “data1” on the storage pool “datasets”:

```
$ cvmkfile -k datasets 2g data1
```

Initializing a Volume (cvmkfs)

Use the `cvmkfs` command to initialize an Xsan volume based on the information in the volume's configuration (in `/Library/Filesystems/Xsan/config/volume.cfg`).

WARNING: Initializing a volume destroys all existing data on the volume.

```
$ sudo cvmkfs [-G] [-F] [volume]
```

Parameter	Description
-G	Don't display "Press return to continue" prompts
-F	Don't display warning and verification prompts. Use this parameter with caution.
<i>volume</i>	The name of the volume to initialize. This name matches the name of a configuration (.cfg) file in <code>/Library/Filesystems/Xsan/config/</code> .

Applying Volume Configuration Changes (cvupdatefs)

Use the `cvupdatefs` command to apply configuration file changes to a volume after you modify the volume's configuration files.

```
$ sudo cvupdatefs [-f] volume [configdir]
```

Parameter	Description
-f	Update without prompting for confirmation or advising of errors in the configuration file.
<i>volume</i>	The volume to update. If you don't specify a volume, available volumes are listed for you to choose from.
<i>configdir</i>	Location of the volume's configuration (.cfg) file if it is not in the default location (<code>/Library/Filesystems/Xsan/config/</code>).

Defragmenting a File, Folder, or Volume (snfsdefrag)

Use the `snfsdefrag` command to defragment a file by reallocating its data in a single extent. This can improve read and write performance for a file by increasing disk efficiency and reducing file metadata management overhead.

To defragment a file or folder:

```
$ snfsdefrag [-D] [-d] [-q] [-s] [-v] [-K affinity]
             [-k affinity] [-m count] [-r] target
```

To report file extents without defragmenting:

```
$ snfsdefrag -e [-K affinity] [-r] target [target] [...]
```

To display an extent count without defragmenting:

```
$ snfsdefrag -c [-K affinity] [-r] target [target] [...]
```

To prune a file (remove allocated extents beyond the end of file):

```
$ snfsdefrag -p [-D] [-v] [-q] [-K affinity] [-m count]
             [-r] target [target] [...]
```


To list files that are candidates for defragmentation:

```
$ snfsdefrag -l [-D] [-v] [-K affinity] [-m count] [-r]
    target [target] [...]
```

Parameter	Description
-c	Display an extent count but don't defragment target.
-D	Display debugging messages.
-d	Operate on files with other than the current depth.
-e	Report extents without defragmenting.
-K <i>affinity</i>	Only operate on files with the specified storage pool affinity. <i>affinity</i> – the affinity key (in Xsan, the affinity key is the same as the name of the storage pool) You can use the <code>cvadmin show long</code> command to see a storage pool's affinity key.
-k <i>affinity</i>	Allocate new extents on the storage pool with this affinity.
-l	List files that might benefit from defragmentation.
-m <i>count</i>	Only operate on files with more than <i>count</i> extents.
-p	Prune instead of defragment.
-q	Suppress messages.
-r [<i>target</i>]	Operate recursively to defragment all files in all folders within the specified target folder.
-s	Allocate new extents on block boundaries.
-v	Display all available information and status during defragmentation.

Examples

Count the extents in the file *datafile*:

```
$ snfsdefrag -c datafile
```

List the extents:

```
$ snfsdefrag -e datafile
```

Defragment the file *datafile*:

```
$ snfsdefrag datafile
```

Defragment every file in the folder */datafolder/* (or any folder within */datafolder/*) that has more than one extent:

```
$ snfsdefrag -r datafolder
```

Recover unused preallocated disk space assigned to every file in folder */datafolder/*:

```
$ snfsdefrag -rp datafolder
```

Controlling the Xsan File System (xsanctl)

Use the `xsanctl` command to control basic Xsan file system functions. For details, see the `xsanctl man` page.

```
$ sudo xsanctl command
```

xsanctl Commands

xsanctl command	Description
<code>ping</code>	Sends a “ping” message to the Xsan file system to verify that it’s responding to management requests.
<code>mount <i>volume</i></code>	Mount an Xsan volume on the computer. <i>volume</i> – the name of the volume
<code>unmount <i>volume</i></code>	Unmount an Xsan volume on the computer. <i>volume</i> – the name of the volume
<code>sanConfigChanged</code>	Notifies the Xsan file system that it should reload the SAN configuration.
<code>disksChanged</code>	Notifies the Xsan file system that it should rescan disks.

Mounting an Xsan Volume

Use the `xsanctl` command to mount an Xsan volume on a computer.

- 1 Either go to the computer and open Terminal, or use SSH to log in to the computer remotely:

```
$ ssh user@computer
```

where *user* is a user account on the remote computer and *computer* is its IP address or DNS name.

- 2 Mount the volume:

```
$ sudo xsanctl mount volume
```

For example:

```
$ sudo xsanctl mount SanVol
```

Unmounting an Xsan Volume

Use the `xsanctl` command to unmount an Xsan volume on a computer.

- 1 Either go to the computer and open Terminal, or use SSH to log in to the computer remotely:

```
$ ssh user@computer
```

- 2 Unmount the volume:

```
$ sudo xsanctl unmount volume
```

For example:

```
$ sudo xsanctl unmount SanVol
```

Viewing Logs

The system log to which Xsan writes information about SANs is in `/var/log/system.log`.

Volume logs are in `/Library/Filesystems/Xsan/data/volume/log/cvlog`, where *volume* is the name of the specific volume.

Xsan Configuration Files

Xsan stores its configuration information in the following files:

Note: None of these files is intended to be edited by Xsan users, except under the direction of an Apple support engineer.

File or folder in <code>/Library/Filesystems/Xsan/config/</code>	Contents
<code>volume.cfg</code>	Volume settings
<code>volume-auxdata.plist</code>	Additional volume settings used by Xsan Admin
<code>fsmlist</code>	Volume autostart list
<code>fsnameservers</code>	Controller list
<code>automount.plist</code>	Xsan volumes to be mounted during startup, and their mount options
<code>config.plist</code>	Private Xsan Admin configuration information
<code>notifications.plist</code>	Notification settings made with Xsan Admin
<code>notes/</code>	Note files whose contents were entered in Xsan Admin's Inspector window
<code>uuid</code>	Private Xsan Admin computer identification information

Examples

Sample configuration files are in `Library/Filesystems/Xsan/examples/`.

affinity A relationship between a folder on an Xsan volume and one or more storage pools that provide storage for the volume. The affinity guarantees that files placed in the folder are stored only on the associated storage pools. Storage pools can differ in capacity and performance, and affinities can be used to assure that data such as video, which requires high transfer speed, is stored on the fastest storage devices.

allocation strategy In Xsan, the order in which data is written to the storage pools that make up a volume. Applicable only if there's more than one storage pool in a volume, and only if the pools are of the same class. Can be **fill**, **round robin**, or **balance**.

balance An Xsan storage pool allocation strategy. Before allocating space on a volume consisting of more than one storage pool, Xsan checks available storage on all pools, and then uses the one with the most free space.

block allocation size An Xsan volume property. The smallest number of bytes that can be reserved on, written to, or read from an Xsan volume.

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

controller In an Xsan storage area network, short for metadata controller. In RAID systems, controller refers to hardware that manages the reading and writing of data. By segmenting and writing or reading data on multiple drives simultaneously, the RAID controller achieves fast and highly efficient storage and access. See also **metadata controller**.

failover In Xsan, the automatic process by which a standby metadata controller becomes the active metadata controller if the primary controller fails.

failover priority On a SAN with more than one controller, specifies which standby controller to try first during failover.

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

file system server See **FSS**.

fill An Xsan storage pool allocation strategy. In a volume consisting of more than one storage pool, Xsan fills up the first pool before writing to the next.

format (verb) In general, to prepare a disk for use by a particular file system.

FSS File system server. The StorNext File System term for the computer that manages metadata in a storage area network (SAN). In Xsan, this is called a metadata controller.

initialize To prepare a disk for use by a particular file system. In Xsan, to prepare a RAID array for use in a storage pool.

label (noun) In Xsan, an identifying name for a LUN. You can assign a label to a LUN before or during setup of an Xsan storage pool.

label (verb) Used by some sources (such as ADIC) to refer to the process of preparing a logical disk for use with a file system. In Xsan, however, *initialize* is used to refer to preparing a disk for use in a storage pool.

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

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

metadata Information about a file system and the files it stores (for example, which disk blocks a file occupies or which blocks are available for use). In Xsan, metadata is managed by a metadata controller and exchanged over an Ethernet connection, while actual file data is transferred over a Fibre Channel connection.

metadata controller The computer that manages metadata in an Xsan storage area network.

mount (verb) To make a remote directory or volume available for access on a local system. In Xsan, to cause an Xsan volume to appear on a client's desktop, just like a local disk.

physical disk An actual, mechanical disk. Compare with **logical disk**.

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

RAID 0 A RAID scheme in which data is distributed evenly in stripes across an array of drives. RAID 0 increases the speed of data transfer, but provides no data protection.

RAID 0+1 A combination of RAID 0 and RAID 1. This RAID scheme is created by striping data across multiple pairs of mirrored drives.

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

RAID 3 A RAID scheme that stripes data across two or more drives and stores parity data on a dedicated drive. In the event of a disk failure, the redundant parity bits can be used to reconstruct data on any drive.

RAID 5 A RAID scheme that distributes both data and parity information across an array of drives one block at a time, with each drive operating independently. This enables maximum read performance when accessing large files.

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

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

RAID set See **RAID array**.

round robin An Xsan storage pool allocation strategy. In a volume consisting of more than one storage pool, Xsan allocates space for successive writes to each available pool in turn.

SAN Storage area network. In general, a network whose primary purpose is the transfer of data between computer systems and storage elements and among storage elements. In Xsan, a SAN is a combination of one or more controllers, storage volumes, and storage clients.

storage pool A group of logical disks that share common characteristics, such as throughput or latency, across which user data is striped. In Xsan, storage pools are combined into volumes. ADIC's StorNext File System calls this a stripe group.

stripe (verb) To write data to successive stripes in a RAID array or LUN.

stripe breadth An Xsan storage pool property. The number of bytes of data, expressed as a number of file system blocks, that Xsan writes to a LUN in a storage pool before moving to the next LUN in the pool.

stripe group The StorNext File System term for an Xsan storage pool.

volume A mountable allocation of storage that behaves, from the client's perspective, like a local hard disk, hard disk partition, or network volume. In Xsan, a volume consists of one or more storage pools. See also **logical disk**.

Index

A

access
 permissions for folders 80
 read-only volumes 82
 unmounting a volume 81
 access control list. *See* ACLs
 ACLs
 enabling and disabling 63
 setting up 80
 Active Directory 40, 44
 adding clients
 to Xsan SAN 72
 adding storage 20, 53
 affinity
 assigning to folder 51, 59
 described 17, 19
 removing 60
 allocation strategy
 setting for volume 63
 availability considerations 30

B

balance
 volume allocation strategy 63
 block allocation size
 choosing for a volume 62

C

clients
 adding 72
 checking quotas from 89
 defined 15
 limit per SAN 21
 moving 73, 74
 removing from SAN 83
 StorNext 117
 using a volume 107
 client worker threads 76
 commands. *See* shell commands
 compatibility

 with other versions of Xsan 10
 with StorNext software 116
 configuration files 131
 controllers
 adding 94
 changing IP address 97
 limit per SAN 21
 listing hosted volumes 97
 overview 15
 cvadmin command 121
 cvaffinity command 124
 cvcp command 124
 cvfsck command 125
 checking volumes 67, 68
 repairing volumes 68
 cvlabel command 126
 cvmkdir command 127
 cvmkfile command 127
 cvmkfs command 128
 cvupdatefs command 128

D

defragmenting volumes 67
 delay access time updates 76
 Directory Access 44
 directory cache size 76
 directory services 26, 40, 44
 Directory Utility 44

E

email notifications 105
 Ethernet
 configuring 42
 guidelines 25
 expanding storage 20, 53

F

failover 95
 forcing 95
 failover priority 95
 Fibre Channel

- configuration requirements 24
- monitoring connection failures 108
- supported switches 24
- files
 - limit per volume 21
 - maximum size 21
 - name length limit 21
- file systems. *See* volumes
- fill
 - volume allocation strategy 63
- firewall, and Xsan Admin 41
- fragmentation 67
- free space
 - checking quota use 87
 - checking storage pool 104
 - checking volume 103

G

- grace period (quota) 86
- graphs
 - controller overall CPU use 104
 - controller overall IP network use 104
- group ID. *See* GID
- groups
 - configuring 43
 - setting up 40

H

- hard quota
 - checking 87
 - defined 86
 - setting 85
- home folders
 - creating local 26, 89

I

- installing Xsan software 45
- IP addresses
 - changing for controller 97

J

- journal
 - choosing location 33
 - described 19

L

- logical unit number. *See* LUNs
- logs
 - controlling number of messages 41
 - viewing 106
- LUNs (logical unit numbers)
 - actual size vs. used size 55, 110
 - adding to existing storage pool 55

- description 16
- limit per storage pool 21
- limit per volume 21
- maximum size 21
- name length limit 21
- overview 15
- preparing 39
- size adjusted downward 110

M

- mail service
 - for notifications 27
- man pages
 - for shell commands 120
- memory
 - requirements 23
- metadata
 - choosing location 33
 - described 19
 - estimating space requirement 34
- mounting a volume
 - from the command line 130
- mount options 76
- moving clients
 - to Xsan SAN 73, 74

N

- naming limits 21
- networks
 - overview 15
- notation
 - conventions for commands 120
- notifications
 - mail service required 27
 - setting up 51, 105

O

- Open Directory 40, 44

P

- pager notifications 105
- permissions
 - user access to folders 80

Q

- quotas
 - checking from client 89
 - checking from command line 88
 - checking usage in Xsan Admin 87
 - described 86
 - example 87
 - grace period 86
 - setting 51, 85

R

- RAID schemes for LUNs 30
- read-only volumes 82
- refresh interval, Xsan Admin 41
- repairing a volume 68
- round robin
 - volume allocation strategy 63

S

- SAN (storage area network)
 - adding 52
 - adding clients 72
 - adding storage 53
 - managing multiple 52
 - moving clients 73, 74
 - name length limit 21
 - removing 52
 - renaming 51
- security considerations 20, 30
- serial number
 - sending expiration notifications 105
- Server Assistant 40
- shared secret file 117
- shell commands
 - cvadmin 121
 - cvaffinity 124
 - cvcp 124
 - cvfsck 125
 - cvlablel 126
 - cvmkdir 127
 - cvmkfile 127
 - cvmkfs 128
 - cvupdatefs 128
 - installed location 119
 - man pages 120
 - snfsdefrag 128
 - xsanctl 130
- snfsdefrag command 128
 - defragmenting files 67
- soft quota
 - checking 87
 - defined 86
 - setting 85
- Spotlight
 - enabling on volume 50
- storage, expanding 53
- storage area network. *See* SAN
- storage pools
 - adding to existing volume 57
 - advanced settings 65
 - checking free space 104
 - described 17
 - limit per volume 21
 - name length limit 21

- reserved names 111
- StorNext File System 115
- stripe breadth
 - and block allocation size 66
 - choosing a value 66
- stripe groups. *See* storage pools
- striping, across LUNs 17, 19
- system requirements 23

T

- time server 43
- troubleshooting
 - can't access RAID system 110
 - can't connect to SAN 109
 - can't install software 109
 - can't mount volume 110
 - client unable to reconnect 112
 - computers not listed 109
 - LUN size adjusted downward 110
 - poor Fibre Channel performance 111
 - reserved storage pool names 111
 - unable to rename volume 110
 - wrong creation date in Mac OS 9 110

U

- unmounting a volume 81
 - from the command line 130
- user ID. *See* UID
- users
 - configuring 43
 - finding 41
 - home folders 40
 - setting up 40

V

- volume configuration file 131
- volumes
 - adding to existing SAN 58
 - checking free space 103
 - checking integrity 68
 - configuration file 131
 - defragmenting 67
 - described 18
 - destroying 69
 - fragmentation 67
 - identifying controller 96
 - listed by controller 97
 - mounting from command line 130
 - name length limit 21
 - repairing 68
 - show clients using 107
 - trouble mounting 110
 - unmounting 81, 130

X

Xsan Admin

- installing 51
- remote SAN management 52

Xsan Admin application

- and firewalls 41
- overview 41
- preferences 41

- refresh interval 41

`xsanctl` command 130

Xsan software

- compatibility with StorNext software 116
- installing 45
- removing (uninstalling) 83
- version compatibility 10

Xsan User Quotas application 89