

IBM TotalStorage SAN File System
(based on IBM Storage Tank™ technology)



Installation and Configuration Guide

Version 2 Release 1

IBM TotalStorage SAN File System
(based on IBM Storage Tank™ technology)



Installation and Configuration Guide

Version 2 Release 1

Note

Before using this information and the product it supports, read the information in "Notices."

Second Edition (June 2004)

This edition applies to the IBM TotalStorage SAN File System and to all subsequent releases and modifications until otherwise indicated in new editions.

Order publications through your IBM representative or the IBM branch office servicing your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for reader's comments is provided at the back of this publication. If the form has been removed, you may address your comments to:

International Business Machines Corporation
Design & Information Development
Department CGFA
PO Box 12195
Research Triangle Park, NC 27709-9990
U.S.A.

You can also submit comments by selecting Feedback at www.ibm.com/storage/support/.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 2003, 2004. All rights reserved.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

About this guide	v
Who should use this guide	v
Notices in this guide	v
Publications	vi
SAN File System publications	vi
SAN File System related publications	vii
Web sites	vii

Chapter 1. Introduction to SAN File System

System	1
What is IBM TotalStorage SAN File System?.	1
What are the major features?	1
What's in the box?	2
Summary of changes	2
Components	5
Terminology	6
Global namespace	6
Filesets	8
Storage management	12
Storage pools	14
Volumes	16
Accessing the administrative interfaces	17
Accessing the administrative command-line interface	17
Accessing the SAN File System console	18

Chapter 2. Installing the SAN File System

System	19
Preparing your environment.	23
SAN considerations.	23
Zoning considerations	23
Security considerations	24
Supported storage subsystems	24
Installing the master metadata server.	25
Preparing the engine for installation	26
Installing software for the master metadata server	30
Setting up the master metadata server	40
Installing subordinate metadata servers	44
Preparing the subordinate engine for installation	44
Installing software for the subordinate metadata server	44
Setting up a subordinate metadata server	45
Setting up the cluster	46
Forming the cluster.	46
Validating cluster installation	46
Setting up clients	46
Installing SDD on clients	47
Installing SAN File System on a Windows client	47
Installing SAN File System on an AIX client	49
Installing SAN File System on a Linux client	51
Installing SAN File System on a Solaris client	53
Configuring SAN File System	54
Configuring metadata servers for SNMP traps.	54
Creating storage pools.	55

Configuring filesets.	55
Placement policies	59
Migrating data	65

Chapter 3. Uninstalling SAN File System

System	69
Uninstalling the package repository	69
Uninstalling the metadata server	69
Uninstalling the SAN File System software from a Windows client	69
Uninstalling the SAN File System software from an AIX client	69
Uninstalling the SAN File System software from a Linux client	70
Uninstalling the SAN File System software from a Solaris client	70

Chapter 4. Upgrading the SAN File System from Version 1.1

System	73
Verifying the current state of the SAN File System	77
Upgrading the first metadata server engine.	80
Preparing the first engine for upgrade	80
Installing software on a metadata server.	84
Restore configuration settings	84
Setting up the metadata server	85
Re-enabling the boot drive mirror	85
Setting the upgraded server as the new master	87
Upgrading additional metadata servers	91
Preparing an engine for upgrade	91
Installing software on a metadata server.	92
Setting up a metadata server engine	92
Completing the upgrade process	93
Upgrading SAN File System on a Windows client	93
Windows client upgrade checklist	93
Preparing a Windows client for upgrading	94
Automate client restart on reboot	94
Upgrading SAN File System on an AIX client	95
AIX client upgrade checklist.	95
Preparing an AIX client for upgrading	95
Installing SAN File System on an AIX client	96
Verifying the installation	97
Upgrading the cluster	97
Revising scripts and policies.	98
Modifying metadata server zoning and LUN masking	99

Chapter 5. Upgrading the SAN File System from Version 2.1

System	101
Upgrading the package repository	103
Upgrading metadata server engines	103
Preparing the metadata server for upgrade	104
Upgrading the administrative server package	104
Upgrading the metadata server package	104
Starting the metadata server engine	105

Upgrading SAN File System on a Windows client	105
Upgrading SAN File System on an AIX client	105
Upgrading SAN File System on a Linux client	106
Linux client upgrade checklist	106
Preparing a Linux client for upgrading	106
Upgrading the SAN File System software	107
Upgrading SAN File System on a Solaris client	107
Solaris client upgrade checklist	108
Preparing a Solaris client for upgrading	108
Committing the upgrade	109

Chapter 6. Backing up the SAN File System 111

Managing backups	111
Backing up using the LUN method	111
Backing up using the file-based (API) method	112
Saving additional SAN File System configuration files	114
One-button data collection	114
Saving a FlashCopy image of a fileset	115

Chapter 7. SAN File System installation commands 117

setupsfs	117
setupstclient	118
tmvt	120

Appendix A. Accessibility 123

Appendix B. Getting help, service, and information 125

Getting help online	125
Getting help by telephone	125

Appendix C. Purchasing additional services. 127

Appendix D. Disaster recovery 129

Appendix E. Notices 131

Trademarks	132
------------	-----

Index 135

About this guide

This topic describes the information that is contained in the Installation and Configuration Guide.

This guide provides information useful to planning, installing and configuring IBM TotalStorage SAN File System.

Who should use this guide

This topic describes the audience for the SAN File System Installation and Configuration Guide.

This guide is intended for people who will install and configure SAN File System hardware and software. Those who install and configure software should have experience and skills in the following areas:

- Networking and network management
- Management of attached storage
- SAN management
- Critical business issues, such as backup, disaster recovery, and security

The installer of SAN File System software should meet the following requirements:

- Knowledge and training in the technology of SAN File System and its functions
- Familiarity with the hardware on which the SAN File System will be installed
- Awareness of the procedures in this document
- Awareness of related installation and service publications

Notices in this guide

This topic describes the notices in the Installation and Configuration Guide.

The following notices are contained with the this guide and convey these specific meanings:

Note: These notices provide important tips, guidance, or advice.

Attention: These notices indicate possible damage to programs, devices, or data. An attention notice appears before the instruction or situation in which damage could occur.

CAUTION:

These notices indicate situations that can be potentially hazardous to you. A caution notice appears before the description of a potentially hazardous procedure step or situation.

DANGER

These notices indicate situations that can be potentially lethal or extremely hazardous to you. A danger notice appears before a description of a potentially lethal or extremely hazardous procedure step or situation.

Publications

This topic describes the publications in the SAN File System library and in related libraries.

SAN File System publications

This topic describes the publications in the SAN File System library.

The following publications are available in the SAN File System library. They are provided in softcopy on the *IBM TotalStorage SAN File System Publications CD* and at www.ibm.com/storage/support. To use the CD, insert it in the CD-ROM drive. If the CD does not launch automatically, follow the instructions on the CD label.

Note: The softcopy version of these publications are accessibility-enabled for the IBM® Home Page Reader.

- *IBM TotalStorage SAN File System Release Notes*

This document provides any changes that were not available at the time the publications were produced. This document is available only from the technical support Web site: www.ibm.com/storage/support

- *IBM TotalStorage SAN File System Software License Information*

This publication provides multilingual information regarding the software license for IBM TotalStorage SAN File System Software.

- *IBM TotalStorage SAN File System Administrator's Guide and Reference, GA27-4317*

This publication introduces the concept of SAN File System, and provides instructions for configuring, managing, and monitoring the system using the SAN File System console and administrative command-line interfaces. This book also contains a commands reference for tasks that can be performed at the administrative command-line interface or the command window on the client machines..

- *IBM TotalStorage SAN File System Basic Configuration for a Quick Start, GX27-4058*

The document walks you through basic SAN File System configuration and specific tasks that exercise basic SAN File System functions. It assumes that the physical configuration and software setup have already been completed.

- *IBM TotalStorage SAN File System Maintenance and Problem Determination Guide, GA27-4318*

This publication provides instructions for adding and replacing hardware components, monitoring and troubleshooting the system, and resolving hardware and software problems.

Note: This document is intended only for trained support personnel.

- *IBM TotalStorage SAN File System Installation and Configuration Guide, GA27-4316*

This publication provides detailed procedures to set up and cable the hardware, install and upgrade the SAN File System software, perform the minimum required configuration, and migrate existing data.

- *IBM TotalStorage SAN File System Messages Reference, GC30-4076*

This publication contains message description and resolution information for errors that can occur in the SAN File System software.

- *IBM TotalStorage SAN File System Planning Guide, GA27-4344*

This publication provides detailed procedures to plan the installation and configuration of SAN File System.

- *IBM TotalStorage SAN File System System Management API Guide and Reference, GA27-4315*

This publication contains guide and reference information for using the CIM Proxy API, including common and SAN File System-specific information.

Note: This document contains information and procedures intended for only selected IBM Business Partners. Contact your IBM representative before using this publication.

SAN File System related publications

These publications are related to SAN File System.

- *IBM TotalStorage[®] Subsystem Device Driver User's Guide, SC26-7637*

Web sites

This topic discusses any Web sites that offer additional, up-to-date information about SAN File System.

The following Web sites have additional information about SAN File System:

- www.ibm.com/storage/support
- www.ibm.com/storage/software/virtualization/sfs

Chapter 1. Introduction to SAN File System

What is IBM TotalStorage SAN File System?

This topic provides a brief overview of IBM TotalStorage SAN File System.

IBM TotalStorage SAN File System is a storage area network (SAN)-based, scalable, and highly-available file system and storage management solution for file aggregation and concurrent data sharing in an open, multi-platform environment. It uses SAN technology, which allows an enterprise to connect a large number of heterogeneous computers and share a large number of heterogeneous storage devices over a high-performance network.

With SAN File System, heterogeneous clients can access shared data directly from large, high-performance, high-function storage systems, such as IBM TotalStorage Enterprise Storage Server[®] (ESS) and IBM TotalStorage SAN Volume Controller. The SAN File System is built on a Fibre Channel network, and is designed to provide superior I/O performance for data sharing among heterogeneous computers. It also provides growth capability and simplified storage management.

SAN File System differs from conventional distributed file systems in that it uses a data-access model that separates file metadata (information about the files, such as owner, permissions, and the physical file location) from actual file data (contents of the files). The metadata is provided to clients by the metadata servers. Clients communicate with the metadata servers only to get the information they need to locate and access the files. Once they have this information, the SAN File System clients access data directly from storage devices through the clients' own direct connection to the SAN. Direct data access eliminates server bottlenecks and provides the performance necessary for data-intensive applications.

SAN File System presents a single, global namespace to clients where they can create and share data using uniform file names from any client or application. Data consistency and integrity are maintained through SAN File System's management of distributed *locks* and the use of *leases*. SAN File System provides locks that enable file sharing among SAN File System clients, and when necessary, provides locks that allow clients to have exclusive access to files. A lease determines the maximum period of time that a metadata server guarantees the locks that it grants to clients. A client must contact the metadata server before the lease period ends in order to retain its locks.

SAN File System also provides automatic file placement through the use of policies and rules. Based on the rules specified in a centrally-defined and managed policy, SAN File System automatically stores data on devices in *storage pools* that are specifically created to provide the capabilities and performance appropriate for how the data is accessed and used.

What are the major features?

This topic summarizes the major features of SAN File System.

Direct data access by exploitation of SAN technology

SAN File System uses a data-access model that allows client systems to access data directly from storage systems using a high-bandwidth SAN, without interposing servers. Direct data access helps eliminate server bottlenecks and provides the performance necessary for data-intensive applications.

Global namespace

SAN File System presents a single, uniform, global namespace view of all files in the system to all of the clients, without manual, client-by-client configuration by the administrator. A file can be identified using the same path and file name, regardless of the client system from which it is being accessed. The global namespace shared directly by clients also reduces the requirement of data replication. As a result, the productivity of the administrator as well as the users accessing the data is improved.

File sharing

All clients, regardless of operating system or hardware platform, have uniform access to the data stored (under the global namespace) in the system. File metadata, such as last modification time, are presented to users and applications in a form that is compatible with the native file system interface of the platform.

Policy based storage and data management

SAN File System is aimed at simplifying the storage-resource management and reducing the total cost of ownership by the policy-based automatic placement of files on appropriate storage devices. The storage administrator can define storage pools depending on specific application requirements and quality of services, and define rules based on data attributes to store the files at the appropriate storage devices automatically. SAN File System provides the storage administrator with policy-based data management that automates the management of storage resources and the data stored on those resources.

What's in the box?

This topic describes what comes with SAN File System.

SAN File System is shipped with three CD-ROMs:

- *SAN File System Software CD* contains the SAN File System software that runs on the SUSE Linux Enterprise platform on prerequisite IBM xSeries® hardware (or equivalent), called *engines*. It also contains SAN File System client software that provides clients with local access to the global namespace on your SAN. The client software must be installed on client machines.
- *SAN File System Software Publications CD* is a multilingual CD that includes the Information Center, user publications in PDF format, license information, and the Overview educational module.
- *Master Console Kit CD* contains the master console software and documentation.

Summary of changes

This section describes the enhancements made to SAN File System in version 2.1.0.

The following list describes the technical changes and enhancements made to SAN File System for version 2.1.0.

- **New Web address for the user interface** — The Web address that is used to access the SAN File System console has changed to:

`https://IP_address:7979/sfs`

- **Software-only solution** — SAN File System is now a software-only solution that leverages the strengths of the standard SuSE Linux Enterprise Server platform.
- **Heterogeneous SAN environment** — SAN File System now supports a wide variety of SAN configurations, which eases scaling to large numbers of storage devices and clients
- **Storage pool access** — The metadata servers require access to the metadata storage (system storage pool). They should not have access to the user data storage (user storage pools). Conversely, the clients require access to the user storage pool. They should not have access to the system storage pools.
- **Non-ASCII Unicode characters** — SAN File System supports both uppercase and lowercase non-ASCII Unicode characters in file names. SAN File System policies also support non-ASCII Unicode characters.
- **Additional client platforms** — SAN File System supports these additional client platforms:
 - AIX[®] 5.2 (32-bit and 64-bit)
 - Red Hat Enterprise Linux 3.0 Advanced Server
 - Sun Solaris 9 (64-bit)

SAN File System also supports AIX and Solaris clients running in a clustered environment.

- **FIFO file-system objects** — SAN File System supports FIFO file-system objects in the global namespace for UNIX-based clients.
- **Opportunistic locks (oplocks)** — Windows-based clients can create and use opportunistic locks (oplocks). SAN File System supports level 1, 2, batch, and filter locks.
- **Application binary support** — Application binaries for UNIX-based clients use the POSIX definition of three sets of three file modes bits: one set of reach user, group, and other.
- **Additional storage subsystem** — SAN File System supports heterogeneous, simultaneously connected storage subsystems on clients with HBA sharing (subject to limitations of the client platform, drivers, and storage vendors). SAN File System conforms to FCP standards and is designed to work with any FCP-compliant storage subsystems for user data storage, including:
 - EMC Symmetrix
 - Hitachi 9900 Series and 9900V Series
 - IBM FAStT 600T, 700, and 900 running firmware version 8.4 and software version 8.41

Note: Only IBM storage subsystem are supported for the system storage pool. SAN File System supports an unlimited number of LUNs for user data storage. However, the amount of user data storage that you can have in your environment is determined by the amount of storage that is supported by the storage subsystems and the client operating systems.

- **High availability, non-disruptive maintenance, and serviceability** — SAN File System has been improved to provide greater availability of the cluster.
 - **Automatic workload failover (filesets and master role)** — SAN File System supports the non-disruptive, automatic failover of the workload. SAN File

System automatically redistributes the filesets of a failed or manually stopped metadata server and, if necessary, reassigns the master role to another metadata server in the cluster. SAN File System uses automatic workload failover to provide non-disruptive maintenance for the metadata servers

- **Automated failback of workload** — If you statically assigned any filesets to a specific metadata server, SAN File System will automatically assign those filesets back to their statically assigned metadata server after the engine hosting the failed metadata server comes back online.
- **Non-disruptive LUN additions** — SAN File System supports adding new LUNs to an existing system without restarting the system.
- **Non-disruptive fileset movement** — You can manually reassign a fileset to another metadata server without disrupting service to the clients.
- **Master console improvements** — In addition to updated software levels, the master console allows you to configure Windows® and IBM Director and to modify the machine name of the storage engines.
- **Serviceability improvements** — Improvements have been made to various serviceability components, including:
 - Metadata checker support of a lost-and-found directory for unreferenced objects.
 - One button data collection utility that has a consistent interface across all platforms.
 - Performance improvements that reduce the impact of logging and tracing.
 - LUN and volume reporting to allow for easy comparison between the SAN configurations of each client.
 - Target Machine Validation Tool (TMVT) hardware and software verification tool. The TMVT runs before the SAN File System setup script and provides a report of installed levels of hardware and software, as well as the required levels of hardware and software.
 - IBM is providing you an optional script to customize your SuSE Linux Enterprise operating system. This script enables you to remove unnecessary packages from your system and to keep on your system only those packages required to run SAN File System. The script, named `ktl.sh`, can be found on the top level directory of the SAN File System Software CD.
- **Usability improvements** — SAN File System has many usability improvements, including:
 - A Refresh button is on many SAN File System console panels to enable you to update the panel with latest information.
 - The dashboard has been simplified to provide you with a quick summary of the entire system.
 - The message IDs in the SAN File System console logs are linked to the Information Center to help you quickly resolve problems.
 - The helps for messages have been improved.
 - You can filter the SAN File System console logs by message date and severity.
 - You can display statistics for fileset transactions, policy sets, and file metadata.
 - You can add and remove individual privileged clients from the administrative command-line interface and SAN File System console.
 - SAN File System supports additional configuration options for the metadata servers and cluster.
 - You can start the cluster and metadata servers in the offline state.
- **Terminology changes** — There are two major terminology changes in SAN File System:
 - The term *container* has been change to the more intuitive term, *fileset*.
 - The administrative command-line interface tool changed from *tanktool* to *sfscli*.

– – The string *tank* in most references has been replaced with *sfs* or *sanfs*.

Components

The following figure illustrates the major components of SAN File System.

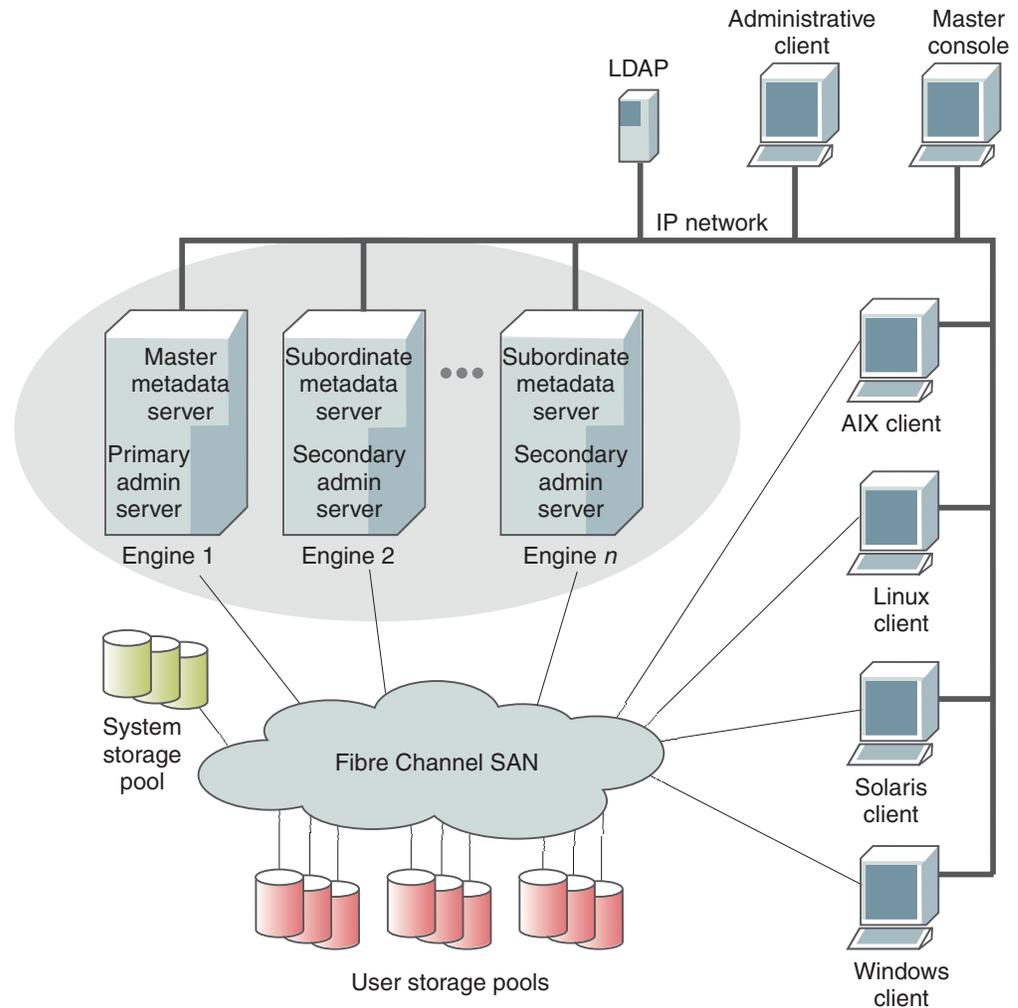


Figure 1. SAN File System components

The metadata servers and clients communicate over a private IP network and access data over a Fibre Channel storage attached network (SAN). SAN File System relies on networking hardware (including an IP network, SAN, network switches, and routers) that already exists in your environment.

The *metadata servers* run on separate physical machines (known as *engines*) and perform metadata, administrative, and storage-management services. The metadata servers are clustered for scalability and availability, and are referred to collectively as the *cluster*. In the cluster, there is one master metadata server and one or more subordinate metadata servers. Additional metadata servers can be added, as required, when the workload grows.

The metadata resides on private storage that is shared among all the metadata servers in the cluster. This storage is known as the *system storage pool*. A storage pool is a collection of SAN File System volumes in the SAN. The system storage

pool contains the system metadata (such as system configuration and state information) and file metadata (such as file creation date and permissions). The actual file data is stored on the *user storage pools*, which may be shared among the clients.

The *administrative server* allows SAN File System to be remotely monitored and controlled through a Web-based user interface, called the *SAN File System console*. In addition, the administrative server processes requests issued from the administrative command-line interface, which can also be accessed remotely. The ability to access the SAN File System through these two types of interfaces allows you to administer SAN File System from almost any system with network connectivity. The administrative server uses an *LDAP server* to look up authentication and authorization information about the administrative users. The primary administrative server runs on the same engine as the master metadata server. It receives all requests issued by administrators and also communicates with the administrative servers that run on each additional metadata server in the cluster to perform routine requests.

Computers that are going to share data and have their storage centrally managed are all connected to the SAN. In SAN File System, these computers are known as *clients*. The SAN File System client software enables the clients to access a single, uniform global namespace through a virtual or installable file system. These clients can act as servers to a broader clientele, providing network File System (NFS) or Common Internet File System (CIFS) access to the global namespace or hosting applications (such as database servers or Web-hosting services that use multiple servers).

The *master console* provides serviceability features, including the remote-support interface (or remote access) and service alert (for call home) capabilities. The master console is a required feature for SAN File System that can be shared with other IBM TotalStorage products, such as SAN Volume Controller.

Terminology

Global namespace

The *global namespace* is the key to SAN File System. It allows common access to all files and directories to all SAN File System clients, and ensures that all SAN File System clients have consistent access and a consistent view of the data and files managed by SAN File System. Having common access to all files reduces the need to store and manage duplicate copies of data and simplifies the backup process. Security mechanisms, such as permissions and access control lists (ACLs), restrict visibility of files and directories.

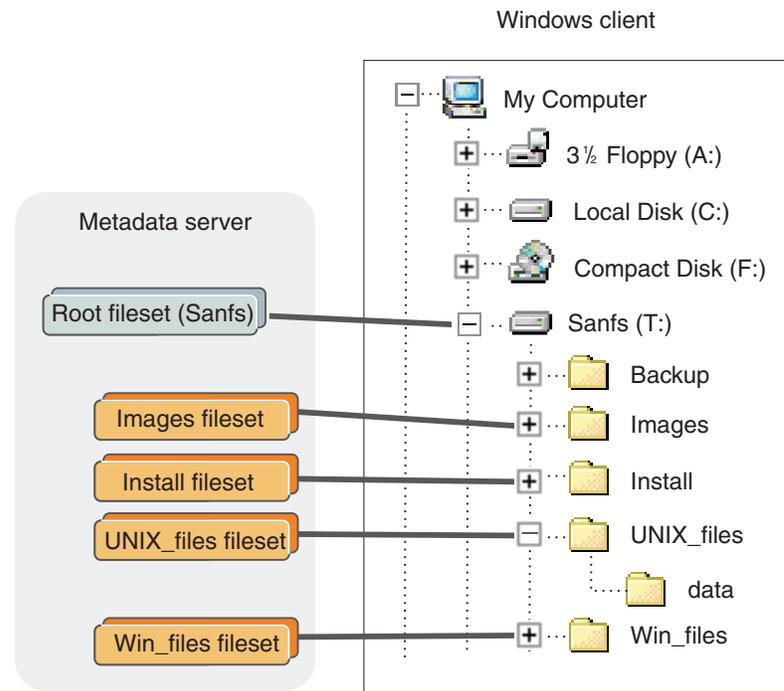
Client access to the global namespace

SAN File System clients mount the global namespace on their systems to access the filesets. After the global namespace is mounted on a client, users and applications can use it just as they do any other file system in order to access data and to create, update, and delete directories and files.

From a client's perspective, the global namespace appears as a normal directory. On a UNIX[®]-based client, the global namespace looks like a mounted file system. On a Windows client, it appears as another drive letter and looks like any other NTFS file system. Basically, the global namespace looks and acts like any other file system on a client's system.

Note: A client cannot move, rename or delete a fileset, and cannot create hard links across fileset boundaries.

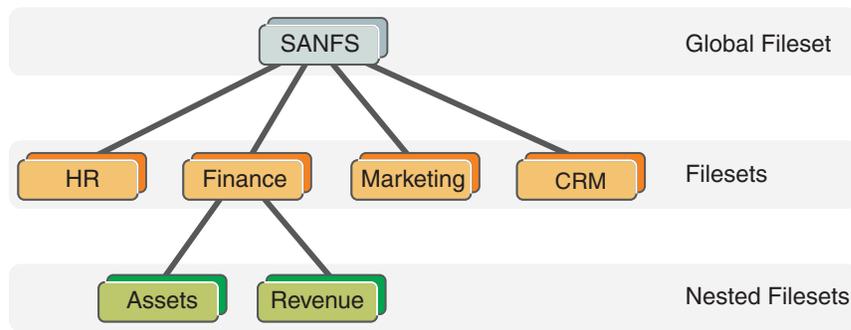
The following figure illustrates the appearance of the fileset from the metadata server and client perspectives. There are five filesets shown: the root, Images, Install, Unix_files, and Win_files. Some of these filesets have subdirectories (for example, the folder Backup is a subdirectory on the root filesystem, and the fileset Unix_files, has a subdirectory named data). The client, however, cannot tell which folders are filesets; they appear all as regular directories.



Global namespace structure

The global namespace is organized into filesets. Each fileset is available to the global namespace at its attachment point. You are responsible for creating filesets and attaching them to directories in the global namespace. This can be done at multiple levels. An attach point appears to a SAN File System client as a directory in which the client can create files and directories (permissions permitting).

The following figure shows a sample global namespace. In this sample, the global fileset is attached to the root level in the namespace hierarchy (sanfs), and the filesets (HR, Finance, Marketing, and CRM) are attached to the global fileset, and the nested filesets (Assets and Revenue) are attached to the Finance fileset. By defining the path of a fileset's attach point, you also automatically define its nesting level in relationship to the other filesets.



Shared access to the global namespace

A *homogeneous environment* is one in which all clients run the same operating system. In a homogeneous environment, SAN File System provides access and semantics that are customized for the operating system that is running on the clients. For example, when files are created and accessed from only Windows clients, all the security features of Windows are available and enforced. When files are created and accessed from only UNIX clients, all the security features of UNIX are available and enforced.

A *heterogeneous environment* is one in which clients run more than one type of operating system. In a heterogeneous environment, there is a restricted form of access. For example, when files created on an UNIX client are accessed by a Windows client, access is controlled using only the semantics and permissions of the “other” permission bits in UNIX. Similarly, when files created on a Windows client are accessed on an AIX client, access is controlled using only the semantics and permissions of the “everyone” group in Windows.

Filesets

In most file systems, a typical file hierarchy is represented as a series of folders or directories that form a tree-like structure. Each folder or directory could contain many other folders or directories, file objects, or other file-system objects, such as symbolic links or hard links. Every file system object has a name associated with it, and it is represented in the namespace as a node of the tree.

SAN File System introduces a new file system object, called a *fileset*. A fileset can be viewed as a portion of the tree-structured hierarchy (or global namespace). It is created to divide the global namespace into a logical, organized structure. Filesets attach to other directories in the hierarchy, ultimately attaching through the hierarchy to the root of the SAN File System cluster mount point. The collection of filesets and their content in SAN File System along with the file system root combine to form the global namespace. Fileset boundaries are not visible to the clients; only the administrator of SAN File System can see them.

From a client’s perspective, a fileset appears as a regular directory or folder within which the clients can create their own regular directories and files. Clients cannot delete or rename the directories that represent filesets.

In addition to organizing the overall structure of the global namespace, SAN File System also uses filesets for these purposes:

- Represent a unit of workload for the metadata servers
- Provide a level of granularity for data replication (using FlashCopy[®] images)

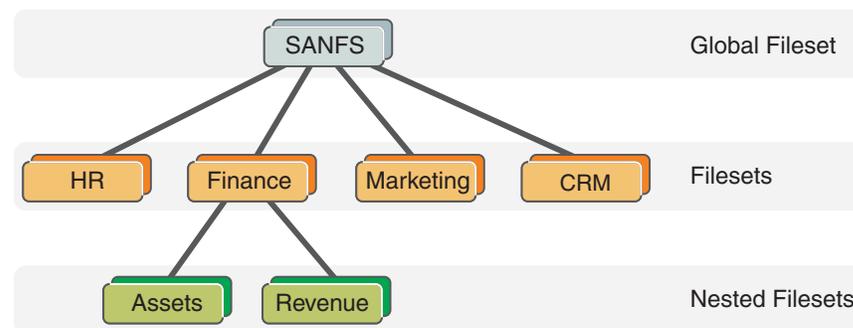
- Control the amount of space used by the clients (through hard and soft quotas)

A fileset has the following properties:

- A fileset name.
- A directory path leading to the directory within which the fileset is attached. The directory path for the global fileset is the same as the cluster name, *sanfs*.
- A directory name that the fileset is given at the end of the directory path.
- A hard or soft quota.

The root of the global namespace is the *global fileset*. The name of the global fileset is always ROOT. The directory path of the global fileset is specified when you set up the global namespace and is the same as the cluster name *sanfs*.

When you create a fileset, you attach it to a specific location in the global namespace, either to the global fileset or to another fileset. When a fileset is attached to another fileset, it is called a *nested fileset*.



You can detach a fileset and reattach it at the same location or a different location. If a fileset is reattached at a different location, all the files contained in the fileset are rooted to the new location without any further operations. Before a fileset can be detached, any nested filesets must be detached first.

Filesets and clients

From a client perspective, a fileset appears to be a regular directory. Users and applications running on the clients can create objects, such as directories and files, within the fileset.

A fileset must be attached to the global namespace before it is available for use by clients.

Users cannot create hard links across fileset boundaries. In addition, a user cannot rename, move, or delete a directory that is the root of a fileset. If a user attempts to perform any of these operations, global namespace issues an error message.

Filesets and metadata servers

When creating a fileset, you may statically assign it to a specific metadata server or SAN File System can dynamically assign it to a metadata server for you. That metadata server is then responsible for providing metadata and locks to clients when they request access to files that reside in that fileset. The fileset-to-metadata server assignment is automatically communicated between clients and metadata servers. The client transparently discovers which metadata server to use when accessing files in a fileset. Each metadata server should be assigned to manage one

or more filesets. If a metadata server is not managing any filesets, it is considered to be in standby mode. You can have an idle, or nearly idle, metadata server available to provide failover, if desired.

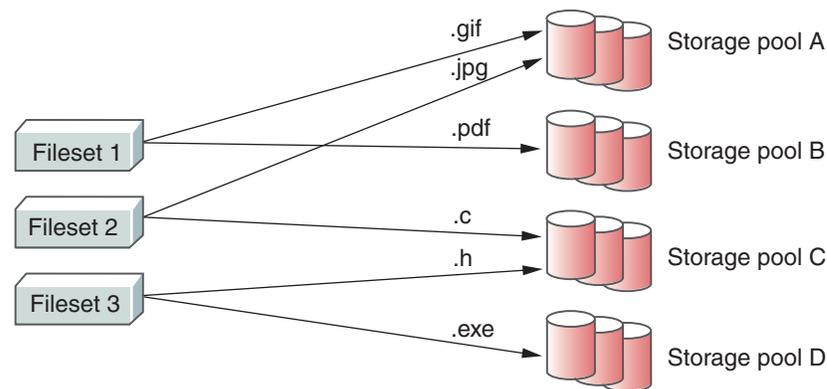
You should create at least one fileset for each metadata server in the cluster. However, creating more filesets gives you greater flexibility in distributing filesets among metadata servers in order to maintain availability and to balance the workload.

Tip: You can assign a nested fileset to a different metadata server than the one to which its parent fileset is assigned.

You can reassign a fileset to another metadata server, for example, to balance the workload. While filesets are being reassigned, they are temporarily unavailable to clients. After the reassignment, the clients can continue transparently and will automatically recognize the new metadata server hosting the fileset.

Filesets and storage pools

Filesets are not specifically related to storage pools, although each file in a fileset physically resides in blocks in a storage pool. This relationship is many-to-many; each file in the fileset can be stored in a different user storage pool. A storage pool can contain files from many filesets. However, all of the data for a particular file is wholly contained within one storage pool. The following illustration shows an example of the relationship between filesets and storage pools.



When you create file placement policies, you can specify that all files created in a particular fileset are to be stored in a specific storage pool.

Fileset considerations

You can create filesets based on conditions in your environment (for example, workflow patterns, security, or backup considerations, all the files used by a specific application, or files associated with a specific application or client). Filesets are used not only for managing the storage space used, but also for creating FlashCopy images. Correctly defined filesets mean that you can take a FlashCopy image for all the files in a fileset together in a single operation, providing a consistent image for all of those files. The global namespace is partitioned into filesets that match the data-management model of the enterprise. Filesets can also be used as a criteria when placing individual files in global namespace.

Tip: When you are creating filesets, consider the overall I/O loads on the cluster. Because each fileset is assigned to one (and only one) metadata server, you need to balance the load across all metadata servers in the cluster by assigning filesets appropriately.

Separate filesets by their *primary allegiance* of the operating system to facilitate file sharing. Separating filesets also facilitates file-based backup methods (for example, utilities, such as tar, and Windows backup applications such as VERITAS NetBackup or IBM Tivoli® Storage Manager); full metadata attributes of Windows files can be backed up from a Windows backup client only and full metadata attributes of UNIX files can be backed up from an UNIX backup client only.

There are no known limitations to the number of filesets that you can create; however, when the number of filesets is greater than one thousand, response time will increase when you issue fileset commands.

Fileset permissions

When you create and attach a new fileset to the global namespace, the fileset is owned by user *Anonymous*. A UNIX root user or a Windows administrator user must change the ownership and permissions of the fileset before the fileset is usable. (You must do this for the FlashCopy directory and the lost+found directory under the fileset root.) You need to make these changes only once in the lifetime of a fileset. The changed permissions are persistent across metadata server restarts and whenever the fileset is detached or attached.

Unlike the requirement for the global fileset, a UNIX or Windows user can own a fileset exclusively. The fileset is not required to have write permissions for both UNIX and Windows domains.

Tip: If you change the permissions of a fileset after you create a FlashCopy image and then revert back to that FlashCopy image, the permissions also revert to their settings at the time when the FlashCopy image was taken.

Fileset quotas

When creating a fileset, you can specify a maximum size for the fileset, called a *quota*, and specify whether SAN File System should generate an alert if the size of the fileset reaches or exceeds a specified percentage of the maximum size, called a *threshold*. For example, if the quota on the fileset is set to 100 GB, and the threshold is 80%, and alert will be generated when the fileset contains 80 GB of data. (Note that the quota is based on space allocated to the fileset, not the data it contains.)

The action taken when the fileset reaches its quota size depends on whether the quota is defined as hard or soft. If a hard quota is used, once the threshold is reached, new client requests to add more space to the fileset (by creating or extending files) are denied. If a soft quota is used, which is the default, more space can be allocated but alerts continue to be sent. Once the amount of physical storage available to global fileset is exceeded, no more space can be used. The quota limit, threshold and quota type can be set individually for each fileset.

Note:

- The space used by a fileset includes the space used by FlashCopy images. It does not include the space used by any filesets nested within it.
- The metadata servers compute and track hard quota limits for filesets in multiples of the partition size. If a hard quota is not set as a multiple of the partition size, quota violation errors appear in the log file even though the size of the fileset has not reached the specified limit. To avoid this problem, specify hard quota limits as multiples of the partition size (for example, if the partition size is 16 MB, set the quota to multiples of 16).

Nested fileset considerations

Consider the following circumstances when creating nested filesets:

- You cannot access a nested fileset if the metadata server that is hosting the parent fileset is unavailable. In other words, if the metadata server hosting a fileset is offline, any nested filesets, even if hosted by a different metadata server, would also be offline.
- A FlashCopy image is created at the individual fileset level and does not include any nested filesets. You cannot make a FlashCopy image of a fileset and any nested filesets in a single operation. This can be of concern if you are required to have a consistent image of a fileset and its nested filesets. Making FlashCopy images in multiple operations could lead to ordering or consistency issues.
- To detach a fileset, you must first detach all of its nested filesets.
- It is not possible to revert to a FlashCopy image when nested filesets exist within the fileset. You must manually detach the nested filesets before reverting to the image. You can reattach the nested filesets after the fileset is reverted.
- When creating nested filesets, attach them only directly to other filesets. Do not attach filesets to client-created directories because a large-scale restore will be more complex.

Storage management

SAN File System provides automatic file placement through the use of policies and storage pools. You can create storage pools that are available to all clients, and define *rules* in policies that cause newly created files to be placed in the appropriate storage pools automatically. For more information about policies, rules, and storage pools, see the related topics below.

File placement

SAN File System provides automatic file placement at the time of creation through the use of policies and storage pools. You can create quality-of-service storage pools that are available to all users and define rules and policies that cause newly created files to be placed in the appropriate storage pool automatically.

A *policy* is a list of rules that determines where the data for specific files is stored. A *rule* is an SQL-like statement that tells a metadata server to place the data for a file in a specific storage pool if the file attribute that the rule specifies meets the specified criteria. A rule can apply to any file being created or to only files being created within a specific fileset depending on how it is defined.

The rules in a policy are processed in order until the condition in one of the rules is met. The data for the file is then stored in the specified storage pool. If none of the conditions specified in the rules of the policy is met, the data for the file is stored in the default storage pool.

Note:

- Rules in a policy are evaluated only when a file is being created. If you switch from one policy to another, the rules in the new policy apply only to newly created files. Activating a new policy does not change the storage pool assignments for existing files. Moving a file does not cause a policy to be applied. You can create multiple policies, but only one policy can be active at a time.
- After a file has been created, you can check its storage pool assignment using the **statfile** command from the administrative command-line interface.
- If you base your policies on user IDs, be aware of how the UNIX **tar** command restores files from backup. During the restore, a file is first created from backup with the user ID of the performer of the backup and

is then changed to the user ID of the original creator of the file. With SAN File System, the policy is applied to the file at creation, so the policy applies to the user ID of the performer of the backup rather than to the user ID of the original file creator.

Policies and rules

This topic describes how SAN File System automates the placement of newly created files into storage pools using policies and rules.

Policies and the rules that they contain are used to automatically assign files to specific storage pools.

Policies

A *policy* is a set of rules that determine where specific files are placed based on the file's attributes. You can define any number of policies, but only one policy can be active at a time. If you switch from one policy to another or make changes to a policy, that action has no effect on existing files in the global namespace. The new or changed policy is effective only on newly created files in SAN File System. Moving a file does not cause the policy to be applied.

A policy can contain any number of rules. There is no limit to the size of a policy.

SAN File System performs error checking for policies in the following phases:

- When you create a new policy, the master metadata server checks the basic syntax of all the rules in the policy.
- When you activate the policy, the master metadata server checks all references to filesets and storage pools. If a rule in the policy refers to a fileset or storage pool that does not exist, the policy is not activated and an error is returned.
- When a new file is created by a client, the rules in the active policy are evaluated in order. If an error is detected, the metadata server responsible for creating the file logs an error, skips all subsequent rules, and assigns the file to the default storage pool. If a default pool does not exist, the file is not created and the metadata server returns an error to the client application.

If your environment is set up in a non-uniform zone configuration (in which clients cannot access all volumes), you need to ensure that the rules in the active policy place files into volumes that are accessible to the clients that use them.

Tip: When SAN File System is first installed, a default policy is created and remains active until you create and activate a new one. The default policy assigns all files to the default storage pool. Although the default storage pool is created when SAN File System is first started, you must assign volumes to it before it can be used. If a user or application on a SAN File System client attempts to create new files that would be assigned to the default storage pool, and there are no volumes assigned to it, the user or application receives No Space errors.

Rules

A *rule* is an SQL-like statement that tells the metadata server to place the data for a file in a specific storage pool if the file meets specific criteria. A rule can apply to any file being created or only to files being created within a specific fileset or group of filesets.

Rules can specify any of these conditions, which when matched, causes that rule to be applied:

- Fileset
- Filename or extension
- Date and time when the file is created
- User ID and Group ID on UNIX clients

SAN File System evaluates rules in order, from top to bottom, as they appear in the active policy. The first rule that matches determines the file's placement. In other words, when a client creates a file, SAN File System scans the list of rules in the active policy to determine which rule applies to the file. When a rule applies to the file, SAN File System stops processing the rules and assigns the file to the appropriate storage pool. If no rule applies, the file is assigned to the default storage pool.

Tip:

- After a file has been created, you can check its storage-pool assignment from the administrative command-line interface using the **statfile** command.
- You can use the **statpolicy** command from the administrative command-line interface to view the statistics about the policy rules.
- If you base your policies on user IDs, be aware of the manner in which the UNIX **tar** command restores files from backup. During the restore, a file is first created from backup with the user ID of the performer of the backup and is then changed to the user ID of the original creator of the file. With SAN File System, the policy is applied to the file at creation, so the policy applies to the user ID of the performer of the backup rather than to the user ID of the original file creator.
- During a restore or migration, a rule that uses the creation date as the placement criteria will assign a file based on the time of the restore or migration, not the original creation time, and a rule that uses a user ID or group ID as the placement criteria will assign a file based on the user ID or group ID of the restore or migration. Therefore, do not use creation time, user ID or group ID to place file.

For detailed information about creating policies and rules, see the related topics below.

Storage pools

A *storage pool* is a named set of SAN File System volumes that can be used to store either metadata or file data. A storage pool consists of one or more volumes that provide a quality of service that you want for a specific use, such as to store all files for a particular application or a specific business division. You must assign one or more volumes to a storage pool before it can be used.

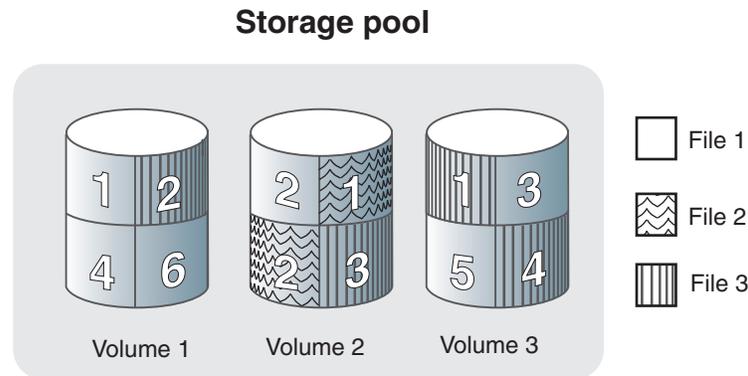
SAN File System has two types of storage pools: system storage pool and user storage pool.

Storage pools and volumes

Typically, you assigns volumes to storage pools based on their common characteristics, such as device capabilities (availability or performance level) and usage (business division, project, application, location, or customer).

Each storage pool manages its own volumes. File space is allocated to the volumes in a given storage pool in a round-robin algorithm (as shown in the following

figure) in logical partitions, or in blocks. Logical partitions are allocated to the system storage pool in 16-MB blocks. For user storage pools, including the default storage pool, you can allocate logical partitions in 16, 64, or 256-MB blocks. All logical partitions in the same storage pool must be the same size.



Tip: You can set a threshold to generate an alert when a storage pool reaches or exceeds a certain percentage of its maximum capacity. By default, an alert is generated when a storage pool becomes 80% full. An alert is logged every five minutes until one or more volumes are assigned to the storage pool. You can set configuration parameters to cause an SNMP trap message to be generated as well. An SNMP trap notifies you of this condition asynchronously.

System storage pool

The *system storage pool* contains the system metadata (system and file attributes, configuration information, and metadata server state) that is accessible to all metadata servers in the cluster. There is only *one* system storage pool that is created automatically when SAN File System is installed. The system storage pool contains the most critical data for SAN File System. The first volume that is assigned to the system storage pool, called the *master volume*, contains the most critical pages of metadata that SAN File System manages.

Important: Use highly-reliable and available logical unit numbers (LUNs) for the system storage pool (for example, mirroring or redundant array of independent disks (RAID), plus hot spares in the backend storage system) so that the cluster always has a robust copy of the system metadata.

Because the amount of metadata grows as the global namespace grows, you must monitor the system storage pool to ensure that there is always enough volumes assigned to it to accommodate the growth. The system storage pool typically requires approximately 2% to 5% of the total storage capacity that SAN File System manages, but this amount varies depending on your environment. Use the alert features on the system storage pool to ensure that you do not run out of space.

Tip: The minimum size of a system volume is 2 GB; therefore, the minimum size of the system storage pool is also 2 GB.

For security and reliability, the volumes that are assigned to the system storage pool should be accessible only to the cluster using a private SAN or a shared SAN with a combination of zoning, LUN masking, or special configuration. For reliability, the volumes should be virtualized RAID arrays (also known as *ranks* within IBM Enterprise Storage Server).

User storage pools

A *user storage pool* contains the blocks of data that make up user files. SAN File System stores the data that describes the files, called file metadata, separately from the actual file data. You can create one or more user storage pools, and then create policies that contain rules that cause metadata servers to store data for specific files in the appropriate storage pools.

The *default storage pool* is a special user storage pool. This optional storage pool is used to store the data for a file if the file is not assigned to a specific storage pool by a rule in the active policy. A default storage pool is created when SAN File System is installed. However, if you want to use the default storage pool, you must assign one or more volumes to it. There can be only one default user storage pool in SAN File System. You can designate any user storage pool that has volumes assigned to it to be the default storage pool. You can choose to disable the default storage pool. In this case, newly created files that do not match any rules in the active policy are not saved.

Volumes

A *logical unit number (LUN)* is the logical unit of storage that a SAN or other disk subsystem can assign to metadata servers and clients. A *volume* is a LUN that is labeled by SAN File System for its use. Volumes are grouped together virtually to form storage pools, in which file data and metadata is stored.

An LUN becomes a SAN File System volume when you add it to a storage pool. It is automatically assigned a system-generated label that identifies it as a SAN File System volume. You must also give the volume a name that is unique among all the volumes used by a SAN File System cluster.

During startup, the metadata server scans all LUNs that it can access in the SAN, searching for the label that tells it that the LUN is a valid SAN File System volume. Clients perform this same search whenever they are started.

System-data LUN operations are performed by the metadata servers. All other data LUN operations are initiated from and coordinated by the metadata servers in the cluster but are actually performed by one or more clients; therefore, the metadata servers no longer need to see the data LUNs, and the clients only need to see the data LUNs that they need to access. This allows SAN File System to support a wide variety of SAN configurations, storage devices, and drivers, and also supports scaling to large numbers of storage devices and clients. This also allows SAN File System to support grouping clients and LUNs into SAN zones to provide enhanced security.

A volume must be empty to be removed from a storage pool. When you remove a volume, SAN File System moves (drains) the contents of that volume across other available volumes in the same storage pool. If the storage pool does not have sufficient space available in other volumes to move all of the data contained in the specified volume, the removal fails and the metadata server suspends the volume (the metadata server cannot allocate new data on that volume).

Tip:

- Keep the storage subsystem device driver's virtual path (vpath) configuration file current. If many LUNs are added and deleted from the metadata server, it is possible for the configuration file to contain references to LUNs that do not exist.

- When the number of entries in the storage subsystem device driver's vpath configuration file reaches 255, any new LUN configured on the metadata server will not be visible.

Volumes and storage pools

When you install SAN File System, there is a system storage pool, which is used by metadata servers to store system and file metadata, and a default storage pool, which can be used to store file data. You can create additional user storage pools for file data; however, no data can be stored in a storage pool until you assign one or more volumes to it. You can also remove the default storage pool if you choose.

The volumes added to the system storage pool are called *system volumes*.

As the amount of metadata that is generated for the server cluster and client files grows, you must ensure that the system storage pool always has enough volumes assigned to it so that it does not run out of space.

You must also ensure that any user storage pools, including the default storage pool, has a sufficient number of volumes. Each storage pool must have at least one volume assigned to it before any files can be stored in it.

To assist you in monitoring storage pool capacity, SAN File System provides a threshold option that you can specify when adding a volume to a storage pool or changing settings for a storage pool. A threshold is a specified percentage of the estimated maximum capacity of the storage pool. When a storage pool reaches or exceeds the percentage specified as its threshold, SAN File System generates an alert. This alert can also generate an SNMP trap message to notify you of the condition asynchronously, if you set the appropriate parameters for SNMP traps.

Limitations to volumes in the system storage pool

The volumes in the system storage pool have these limitations:

- All volumes in the system storage pool must be of the same type of backend storage device and must be one of the supported IBM storage subsystems. You can use IBM TotalStorage SAN Volume Controller to provide mixed storage as long as only the SAN Volume Controller virtual devices are visible to the cluster.
- All volumes in the system storage pool must be visible to all metadata servers in the cluster.
- Each volume in the system storage pool must be at least 2 GB in size.
- The system storage pool is limited to 126 dual-path volumes.

Accessing the administrative interfaces

This section discusses the tasks for accessing the SAN File System console and administrative command-line interface.

Accessing the administrative command-line interface

This topic describes how to bring up the administrative command-line interface.

SAN File System provides a wizard to step you through the process of creating a FlashCopy image.

You must have Administrator, Operator, or Backup privileges to perform this task.

1. Log in directly to the engine, or from another workstation through SSH, using the local operating system authentication mechanism.

2. Log in to the administrative server on the engine using the same administrative user ID and password that you would use to log into the SAN File System console. You can specify the password in one of two ways:
 - Set the password in the `sclif.properties` file, located in your home directory on the engine (for example, `joe/sclif.properties`), to your valid LDAP password using the **tankpasswd** utility.
 - Set the `SFS_CLI_PASSWDFILE` environment variable to the location of the password file.
3. Enter the **sfsccli** command to start the administrative command-line interface, or `sfsccli` session, to run commands in interactive mode.

Accessing the SAN File System console

This topic describes how to bring up the SAN File System console from the master console.

1. Open a supported Web browser window and enter the URL:
`https://master_metadata_server:port/sfs`, where *master_metadata_server* is the name or IP address of the master metadata server, and *port* is the port number of the master metadata server. The default port number is 7979.
2. When the SAN File System signon screen displays, enter your administrator ID and password.

Chapter 2. Installing the SAN File System

This topic provides a high-level overview of the procedures for installing SAN File System in your environment.

Before you begin installing SAN File System, ensure you have access to the planning worksheets. You can find these worksheets (and information about filling them out) in the *SAN File System Planning Guide*, GA27-4344.

You need to also make sure that you have access to all of the required software. Most of the prerequisite software that you need is available on the SAN File System CD-ROM. However, you need to obtain the following software:

- SUSE LINUX Enterprise Server 8.0. You need a licensed copy of LINUX Enterprise Server 8.0 for each of the metadata server engines in the cluster. For more information about obtaining Enterprise Server, visit www.suse.com.
- QLogic driver. For information about obtaining the QLogic driver, visit http://www.qlogic.com/support/oem_detail_all.asp?oemid=22.
- United Linux Service Pack 3. For more information about obtaining the United Linux Service Pack 3, visit www.suse.com.

Checklist

Use the following checklist to install the SAN File System.

Steps	For more information...	
1	Prepare your environment.	"Preparing your environment" on page 23
	a Prepare the SAN.	"SAN considerations" on page 23
	b Prepare switch zoning.	"Zoning considerations" on page 23
	c Prepare for security.	"Security considerations" on page 24
	d Prepare storage devices.	"Supported storage subsystems" on page 24
2	Install the master console.	<i>Master Console User's Guide</i>
3	Install the master metadata server.	"Installing the master metadata server" on page 25
	a Prepare the engine for installation.	"Preparing the engine for installation" on page 26
	1 Cable the metadata server.	"Cabling" on page 26
	2 Obtain software prerequisites.	"Obtain prerequisite software" on page 30
	3 Upgrade system BIOS, if needed.	"Upgrading system BIOS" on page 30
	4 Mirror boot drives.	LSI Logic Configuration Program documentation
	b Install all software.	"Installing software for the master metadata server" on page 30

Steps		For more information...
1	Install SUSE LINUX Enterprise Server 8.0.	"Installing the operating system" on page 31
2	Disable the X Window System.	"Disabling the automatic starting of the X Window System" on page 34
3	Set the date and time on the engine.	"Setting the time and date on the Metadata servers" on page 34
4	Apply United Linux Service Pack 3 (SP3) updates.	"Apply United Linux Service Pack 3 (SP3) updates" on page 34
5	Install prerequisite software.	"Installing prerequisite software" on page 35
	a QLogic driver.	"Install QLogic driver" on page 35
	b MPCLI.	"Install MPCLI" on page 38
	c Java runtime environment.	"Install the Java Runtime Environment" on page 38
	d IBM Director.	"Install the IBM Director agent" on page 38
	e Eclipse.	"Install Eclipse" on page 38
	f Ibmusbasm.	"Install ibmusbasm" on page 38
	g Openslp.	"Install OpenSLP" on page 39
	h IBM SDD.	"Install the IBM Subsystem Device Driver (SDD)" on page 39
	i WebSphere 5.0 Express.	"Install IBM WebSphere 5.0 Express" on page 39
6	Install the SAN File System software.	"Installing SAN File System software" on page 40
c	Set up the master metadata server engine.	"Setting up the master metadata server" on page 40
1	Configure the RSA II adapter.	"Configuring the RSA II adapter" on page 40
2	Upgrade RSA II firmware, if needed.	"Upgrading RSA II firmware" on page 42
3	Configure the metadata server.	"Running the SAN File System setup utility" on page 42
4	Install all subordinate metadata servers.	"Installing subordinate metadata servers" on page 44
	a Prepare the engine for installation.	"Preparing the subordinate engine for installation" on page 44
	1 Cable the metadata server.	"Cabling" on page 26
	2 Obtain software prerequisites.	"Obtain prerequisite software" on page 30
	3 Upgrade system BIOS, if needed.	"Upgrading system BIOS" on page 30
	4 Mirror boot drives.	LSI Logic Configuration Program documentation
	b Install all software.	"Installing software for the subordinate metadata server" on page 44

Steps		For more information...
1	Install SUSE LINUX Enterprise Server 8.0.	"Installing the operating system" on page 31
2	Disable the X Window System.	"Disabling the automatic starting of the X Window System" on page 34
3	Set the date and time on the engine.	"Setting the time and date on the Metadata servers" on page 34
4	Apply United Linux Service Pack 3 (SP3) updates.	"Apply United Linux Service Pack 3 (SP3) updates" on page 34
5	Install prerequisite software.	"Installing prerequisite software" on page 35
	a QLogic driver.	"Install QLogic driver" on page 35
	b MPCLI.	"Install MPCLI" on page 38
	c Java runtime environment.	"Install the Java Runtime Environment" on page 38
	d IBM Director.	"Install the IBM Director agent" on page 38
	e Eclipse.	"Install Eclipse" on page 38
	f Ibmusbasm.	"Install ibmusbasm" on page 38
	g Openslp.	"Install OpenSLP" on page 39
	h IBM SDD.	"Install the IBM Subsystem Device Driver (SDD)" on page 39
	i WebSphere 5.0 Express.	"Install IBM WebSphere 5.0 Express" on page 39
6	Install the SAN File System software.	"Installing SAN File System software" on page 40
c	Set up the subordinate metadata server engine.	"Setting up a subordinate metadata server" on page 45
1	Copy files from the master metadata server.	"Copying tank.properties and the truststore" on page 45
2	Configure the RSA II adapter.	"Configuring the RSA II adapter" on page 40
3	Upgrade RSA II firmware, if needed.	"Upgrading RSA II firmware" on page 42
4	Configure the metadata server.	"Running the SAN File System setup utility" on page 45
5	Set up the cluster.	"Setting up the cluster" on page 46
	a Start the cluster.	"Forming the cluster" on page 46
	b Validate cluster installation.	"Validating cluster installation" on page 46
6	Set up Windows clients.	"Installing SAN File System on a Windows client" on page 47
	a Optionally, install SDD, RDAC, or other multipathing software.	For information about installing SDD, see "Installing SDD on clients" on page 47
	b Obtain Windows client software.	"Obtain version 2.1 software for a Windows client" on page 47

Steps		For more information...	
	c	Install the software.	"Installing the SAN File System software" on page 48
	d	Validate client installation.	"Validating the installation of SAN File System on a Windows client" on page 48
7		Set up AIX clients.	"Installing SAN File System on an AIX client" on page 49
	a	Optionally, install SDD, RDAC, or other multipathing software.	For information about installing SDD, see "Installing SDD on clients" on page 47
	b	Obtain AIX client software.	"Obtain version 2.1 software for an AIX client" on page 49
	c	Install the software.	"Installing the SAN File System software" on page 50
	d	Validate client installation.	"Validating the installation of SAN File System on an AIX client" on page 51
8		Set up Linux clients.	"Installing SAN File System on a Linux client" on page 51
	a	Optionally, install SDD, RDAC, or other multipathing software.	For information about installing SDD, see "Installing SDD on clients" on page 47
	b	Obtain Linux client software.	"Obtain version 2.1 software for a Linux client" on page 51
	c	Install the software.	"Installing the SAN File System software" on page 52
	d	Validate client installation.	"Validating the installation of SAN File System on a Linux client" on page 52
9		Set up Solaris clients.	"Installing SAN File System on a Solaris client" on page 53
	a	Optionally, install SDD, RDAC, or other multipathing software.	For information about installing SDD, see "Installing SDD on clients" on page 47
	b	Obtain Solaris client software.	"Obtain version 2.1 software for a Solaris client" on page 53
	c	Install the software.	"Installing the SAN File System software" on page 53
	d	Validate client installation.	"Validating the installation of SAN File System on a Solaris client" on page 54
10		Perform initial configuration.	"Configuring SAN File System" on page 54
	a	Configure SNMP traps on all metadata servers.	"Configuring metadata servers for SNMP traps" on page 54
	b	Configure storage pools.	"Creating storage pools" on page 55
	c	Configure filesets.	"Configuring filesets" on page 55
	d	Configure placement policies.	"Placement policies" on page 59

Steps		For more information...
e	Migrate data.	"Migrating data" on page 65
1	Estimate the time to migrate.	"Estimating the time to migrate" on page 65
2	Import data.	"Importing data into the SAN File System" on page 66
3	Validate migrated data.	"Verifying the data integrity of migrated data" on page 67
11	Back up the complete system.	Chapter 6, "Backing up the SAN File System," on page 111

Preparing your environment

This topic provides an overview of preparing your environment for the installation of the SAN File System.

Consider the following areas when preparing your environment for the installation of the SAN File System. For additional information about these considerations, refer to the *SAN File System Planning Guide*, which also contains planning worksheets.

- SAN considerations.
- Zoning considerations.
- Security considerations.
- Supported storage subsystems.

SAN considerations

This topic describes the SAN considerations for the SAN File System.

Use the following guidelines to prepare your SAN for the SAN File System.

- Set up your switch configuration to maximize the number of physical LUNs addressable from the Metadata servers and to minimize sharing of fabrics with other non-SAN File System users whose usage may be disruptive to the SAN File System.
- Verify that the storage devices that will be used by SAN File System are set up so that the appropriate storage LUNs are available to the SAN File System.

Zoning considerations

This topic describes the zoning considerations for the SAN File System.

Use the following guidelines to implement zoning for the SAN File System.

Note: For more information about planning to implement zoning, see the *SAN File System Planning Guide*.

- Due to the restriction on the number of LUNs the Metadata servers can currently access, make sure you limit the number of paths created through the fabrics from each metadata server to the storage to two paths, one per host-bus adapter (HBA) port. Some combination of zoning and physical fabric construction may be used to reduce or limit the number of physical paths. Each fabric should consist of one or more switches from the same vendor.

- Keep in mind that there is no level of zoning you can do on a SAN that will protect SAN File System systems from SAN events caused by other non-SAN File System systems connected to the same fabric. Therefore, you should not create fabrics that include traffic and administrative contact from non-SAN File System systems. You can utilize VSANs to accomplish this fabric isolation.
- When metadata storage and user storage reside on the same storage subsystem, you must ensure that the metadata storage is fully isolated and protected from access by client systems. With some subsystems, access to various LUNs is determined by connectivity to various ports of the storage subsystems. With these storage subsystems, hard zoning of the attached switches may be sufficient to ensure isolation of the metadata storage from access by client systems. However, with other storage subsystems (such as ESS), LUN access is available from all ports and LUN masking *must* be used to ensure that the Metadata servers are the only systems allowed to access the metadata storage LUNs.

Note: The SAN File System user LUNs and SAN File System metadata LUNs should not share the same ESS 2105 Host Adapter ports.

- SAN File System clients should be zoned or LUN masked such that each can see user storage only.
- Specify that the Metadata server storage or LUNs are to be configured to the Linux mode (if the metadata storage subsystem has operating system-specific operating modes).

Security considerations

This topic describes the security considerations for the SAN File System.

Verify that Lightweight Directory Access Protocol (LDAP) is set up and configured properly. In addition, you will need to add SAN File System users to the LDAP user database.

Supported storage subsystems

This topic describes the storage subsystems that are supported by the SAN File System.

SAN File System supports heterogeneous, simultaneously-connected Fibre Channel storage subsystems on clients with host bus adapter (HBA) sharing, subject to the limitations of the client platform, drivers, and storage vendors.

SAN File System supports an unlimited number of LUNs for user data storage. However, the amount of user data storage that you can have in your environment is determined by the amount of storage that is supported by the storage subsystems and the client operating systems.

For more information about supported storage subsystems, refer to the following Web site:

www.ibm.com/storage/support/sanfs

System storage pool

Currently, SAN File System supports only these storage subsystems for use in the system storage pool:

- The IBM TotalStorage Enterprise Storage Server (ESS), models 2105-F20 and 2105-800

- The IBM TotalStorage SAN Volume Controller, model 2145 with storage subsystems that are supported by SAN Volume Controller
- IBM FAStT 600T, 700, and 900 running firmware version 8.4 on the storage device and software version 8.41 on the client platforms

Note: Ensure that the IDs for any LUNs that are used by the system storage pool starts with 0. Refer to your storage documentation for information about assigning LUN IDs.

Refer to the IBM storage Web site for the supported code levels of these storage subsystems.

www.ibm.com/storage/support

User storage pool

For user storage pools, SAN File System is designed to work with FCP-compliant storage subsystems that meet the following qualifications:

- Conforms to SCSI standards for device driver interface, including unique device identification.
- Supports the required device drivers and operating-system stack.
- Are SAN-attached to the client machines.

Tip: Consider any restrictions imposed by the storage subsystem, host bus adapter (HBA) cards, device drivers, and client platforms that will be used in your SAN File System environment to ensure that they are all compatible.

Storage subsystems other than ESS or SAN Volume Controller may require additional, manual configuration to be detected and used by SAN File System.

Refer to the platform support documentation for a list of storage subsystems that are supported for each client platform in your environment.

Installing the master metadata server

This topic provides an overview of the steps required to install the SAN File System master metadata server.

Make sure that you have fulfilled the following prerequisites before installing master metadata server:

- You have unpacked the engine and installed (but not configured) the RSA II adapter card
- You have installed the engine in a rack.
- You have attached a keyboard, monitor, and mouse to the engine. Alternatively, a KVM as the console for the engine.

Use the documentation that comes with your engine and with the RSA II adapter to meet these prerequisites.

1. Prepare the engine for installation. This includes these tasks:
 - a. Ensuring that the engine is properly cabled.
 - b. Ensuring that you obtain all prerequisite software.

See “Preparing the engine for installation” on page 26.

2. Install all software, including the operating system, prerequisite software, and the SAN File System software. See “Installing software for the master metadata server” on page 30.
3. Set up the master metadata server to recognize all appropriate storage devices. In addition, you need to configure the RSA II adapter and run the SAN File System setup utility. See “Setting up the master metadata server” on page 40.

Preparing the engine for installation

This topic describes how to prepare an engine to be installed.

1. Ensure that the engine, including the RSA II adapter is cabled properly. See “Cabling.”

Note: If you are using an existing engine that currently has the RSA II adapter installed, disable the RSA II watchdogs to prevent the engine from automatically restarting during the installation process. See “Disabling the RSA II watchdogs” on page 81.

2. Obtain software to be used on the installation process. See “Obtain prerequisite software” on page 30.
3. Update the system BIOS, if necessary.
4. Use the LSI Logic Configuration Program documentation provided with the engine to mirror the boot drive.

Cabling

This topic describes how to cable the metadata server engine and the RSA II adapter.

1. Perform the following steps to cable the hardware:
 - a. For each of the two power cords, connect the appropriate end of the power cord to a power supply and the opposite end to a properly wired and grounded electrical outlet.
 - b. Connect one end of the two Fibre Channel cables to the HBA ports located in expansion slot 2, and connect the opposite end of each cable to the SAN through a switch. See Figure 4 on page 28 for a cabling example.
 - c. For redundancy, connect another Fibre Channel cable to the other HBA port in expansion slot 2, and to the other switch (or zone). This is optional, but recommended.

Note: The SAN File System user LUNs and SAN File System metadata LUNs do not share the same ESS 2105 Host Adapter ports. User LUNs should not be visible to the metadata servers in the SAN File System cluster. Metadata server LUNs should not be visible to SAN File System clients.

- d. Connect one end of the Ethernet cable to the integrated 10/100/1000 Ethernet port in the engine, and connect the opposite end to the Ethernet switch or hub. The example in Figure 4 on page 28 shows a hub.
- e. An Advanced System Management (ASM) connector and USB cable are provided with the RSA II adapter.
 - 1) Connect the USB cable to a USB port on the engine and the other end to the RSA card.
 - 2) With an RJ-45 cable, connect one Ethernet connector on the RSA card to the Ethernet switch or hub that is provided as shown in Figure 4 on page 28.

- 3) Connect a ASM breakout cable (dongle) to the ASM connector on the RSA II card in each of the engines present in the SAN FS cluster. Connect the ASM breakout cable to the previous and next ASM breakout cables with RJ-45 cables. The first and last RJ-45 sockets in the chain must be terminated with the terminators provided. See Figure 2. and Figure 3.
- 4) The 9-pin D-shell serial connector on the ASM connector is not used.

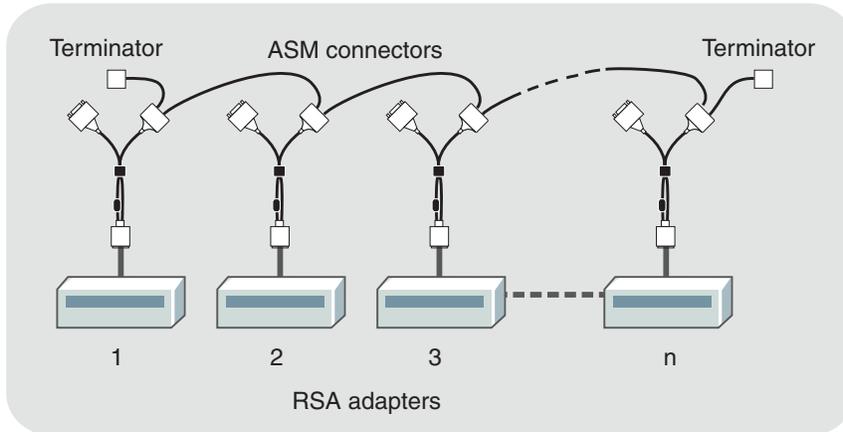


Figure 2. Connecting the RSA adapters together

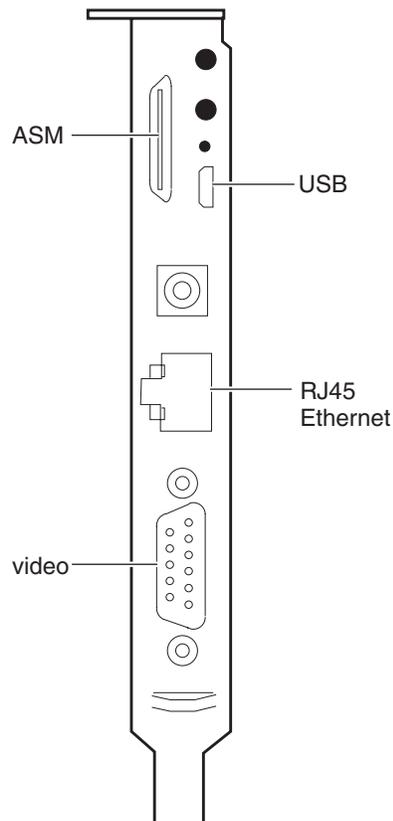


Figure 3. RSA II adapter connectors

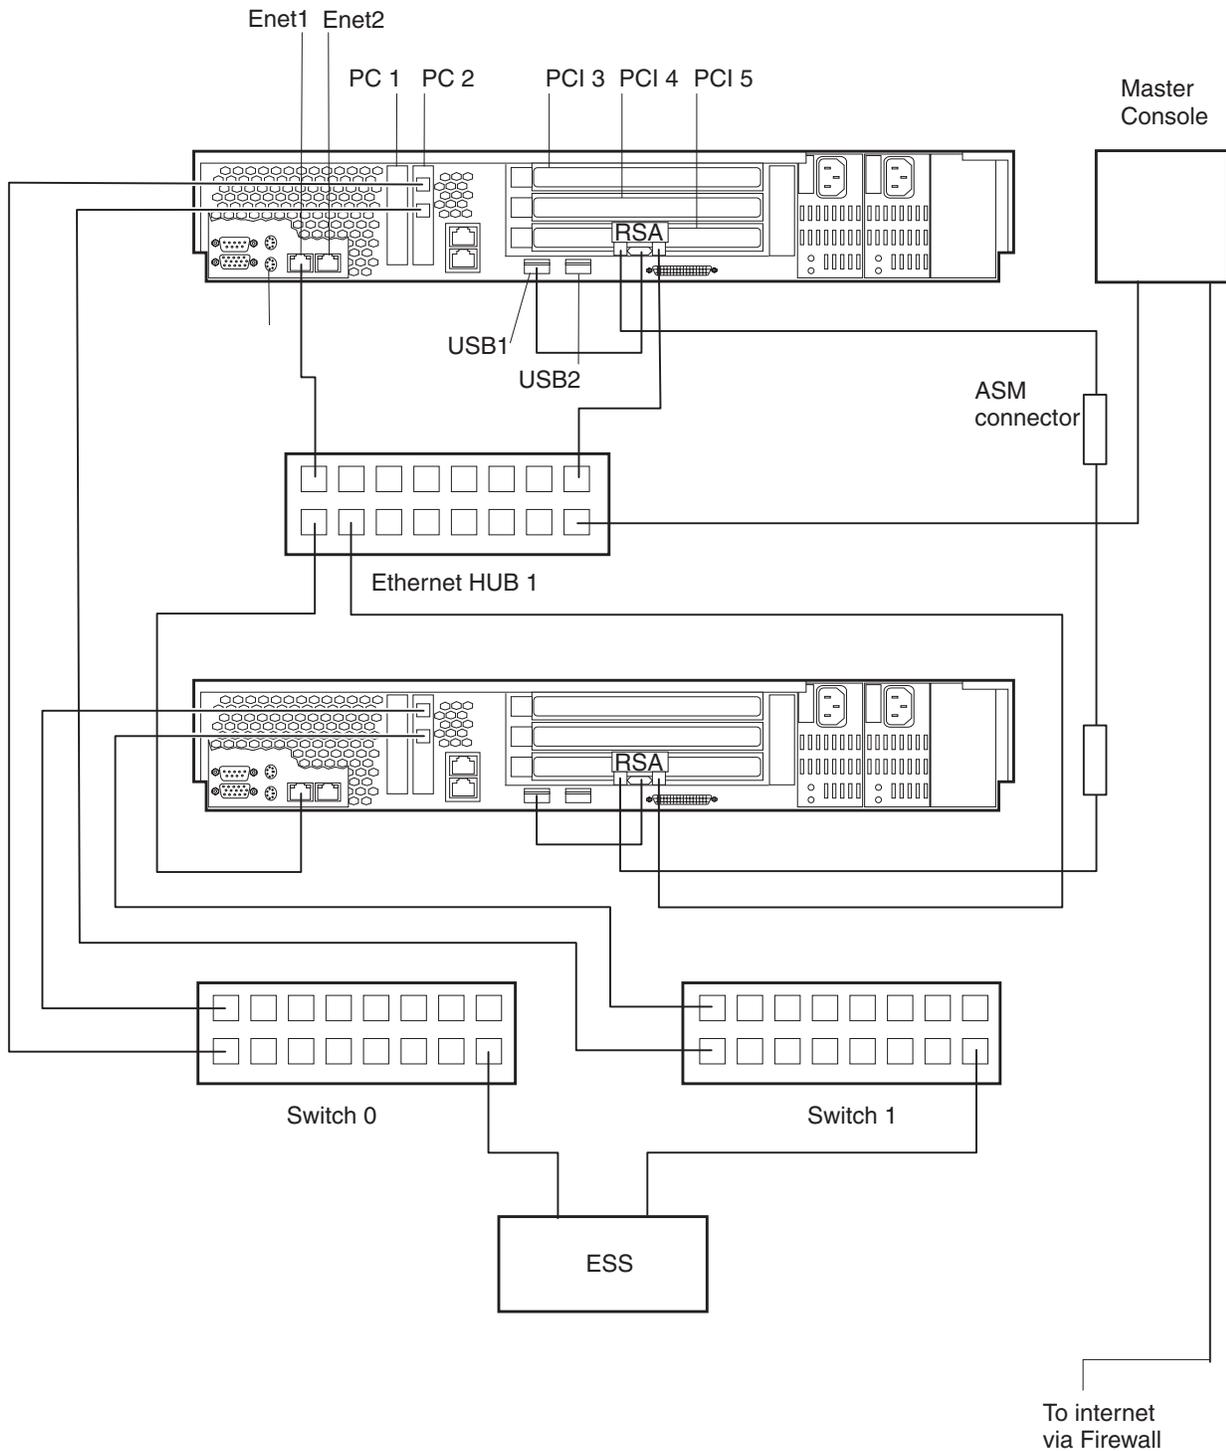


Figure 4. Two-node, two-switch, two-hub cabling example

- f. Use cable clamps to secure the cables across the rear of the engine.
- g. Route the cables along the cable-management-arm channel, securing them with cable straps.

Attention: Interconnect cables to the RS-485 connectors may be too short to route in the cable management arms. Use care when sliding out an engine to avoid damaging a cable or connector.

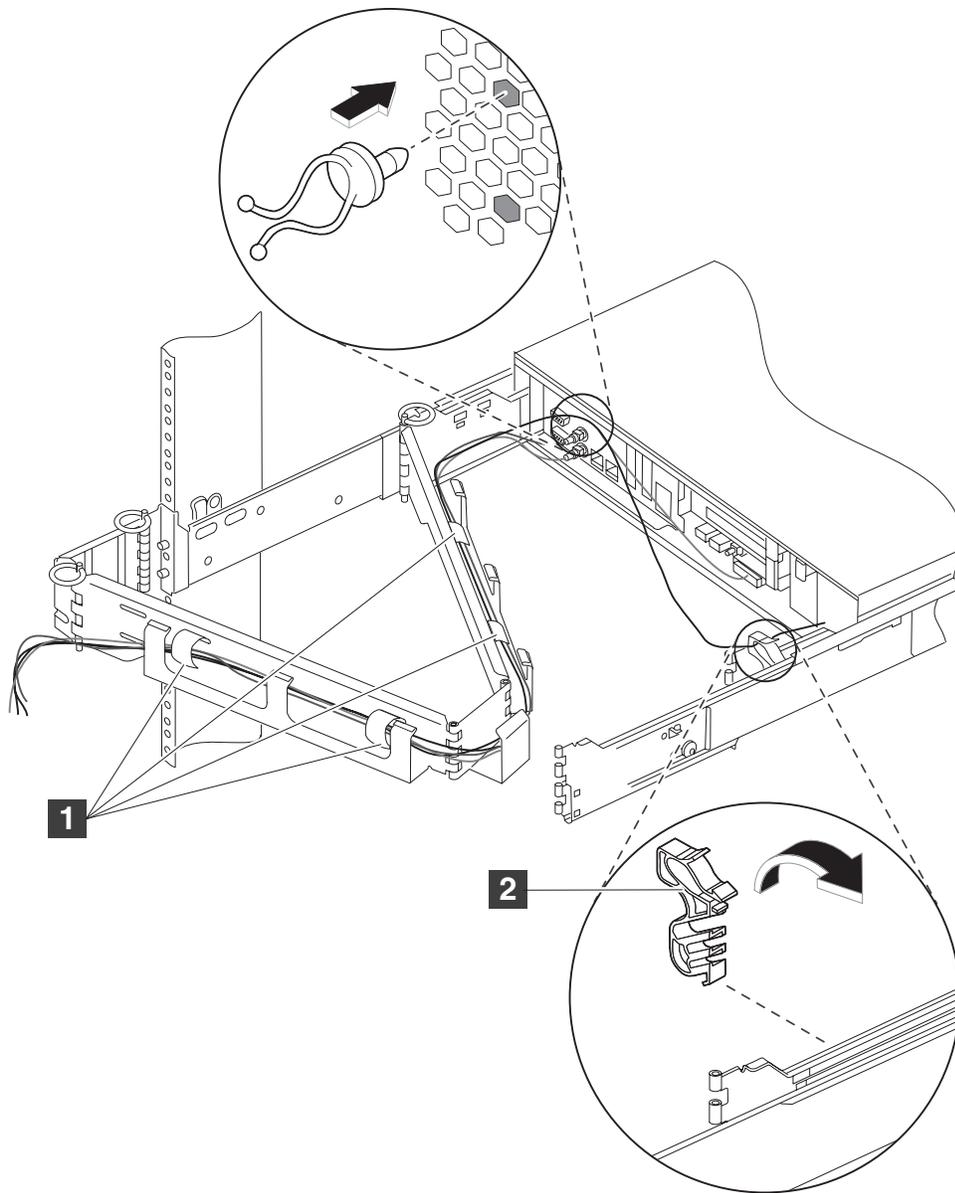


Figure 5. Attaching the cable straps

- 1 Cable straps
 - 2 Cable-restraint bracket
- h. Secure the cable-restraint bracket to the slide rail, if not already done. Route the power and network cables through the cable-restraint bracket, allowing slack in all cables to avoid tension.
 2. For the master console, perform the following steps to cable the hardware:
 - a. Connect one Ethernet adapter port to the internet by way of the corporate firewall.
 - b. Connect the other Ethernet adapter port to the customer's intranet. This network includes the RSA II adapters, Metadata server, and SAN File System clients. The example Figure 4 on page 28 shows an Ethernet hub.
 - c. Attach the keyboard, display, and mouse to the KVM connectors on the master console.

Obtain prerequisite software

Before you begin installing the SAN File System, make sure that you have access to all of the required software. Most of the prerequisite software that you will need is available on the SAN File System CD-ROM.

In addition to the software provided with the SAN File System, you will need to obtain the following software:

- SUSE LINUX Enterprise Server 8.0. You will need a licensed copy of LINUX Enterprise Server 8.0 for each of the metadata server engines in the cluster. For more information about obtaining Enterprise Server, visit www.suse.com.
- QLogic driver (QLA2300F - version 6.06.64). For information about obtaining the QLogic driver, visit http://www.qlogic.com/support/oem_detail_all.asp?oemid=22.
- United Linux Service Pack 3. For more information about obtaining the United Linux Service Pack 3, visit www.suse.com.

Upgrading system BIOS

This topic describes how to upgrade the system BIOS for the metadata server engine.

You need to verify that your system BIOS is at version 1.16 (GEJT56A). To determine the current version of the system BIOS, reboot the engine and watch for the BIOS version to be displayed.

To obtain the correct version of the system BIOS, visit this Web site:

<http://www-307.ibm.com/pc/support/site.wss/MIGR-43902.html>

Note: If you are using an IBM TotalStorage 4146, the BIOS is the same as that used for the IBM eServer™ xSeries 345.

Make sure that you follow the instructions in the README to upgrade the system BIOS. In addition, you will need to copy the BIOS file to a diskette after downloading it. Instructions for downloading and creating the diskette are available on the Web site.

Note: A prerequisite for version 1.16 is the Integrated System Management Processor firmware version 1.05 or later. To obtain this firmware, search for "integrated system management processor" from the IBM Support Web site. You will need a diskette for this firmware update as well.

After updating the BIOS, you may see the following BIOS error messages:

162 configuration error
184 Power Password becomes invalid.

To clear these error messages, press **F1** when prompted during POST to enter the Configuration/Setup Utility. At the Main Menu of Setup, select the option for Save configuration. Then, exit the Configuration/Setup Utility and reboot the system.

Installing software for the master metadata server

This topic provides an overview of the steps that are required to install the software for the master metadata server.

1. Install SUSE LINUX Enterprise Server 8.0. See "Installing the operating system" on page 31

2. Disable the automatic starting of the X Window System (if you choose to have graphical mode as the default desktop setting during the installation of the operating system). See “Disabling the automatic starting of the X Window System” on page 34
3. Set the time and date on the metadata server engine. See “Setting the time and date on the Metadata servers” on page 34
4. Apply the United Linux Service Pack 3. See “Apply United Linux Service Pack 3 (SP3) updates” on page 34.
5. Install all prerequisite software, such as the QLogic driver and the IBM Director agent. See “Installing prerequisite software” on page 35.
6. Install the SAN File System software. See “Installing SAN File System software” on page 40.

Installing the operating system

This topic describes the procedure for installing SUSE LINUX Enterprise Server 8.0 on a metadata server engine.

Important: Use the following procedure to install the SUSE Linux Operating System and then use the procedure in “Apply United Linux Service Pack 3 (SP3) updates” on page 34 to apply the Service Pack. Do not use the United Linux Service Pack 3 CD Kickstart functionality because the installation may not complete successfully.

1. Insert the SLES8 CD-ROM 1 in the CD-ROM drive and reboot the engine.
2. When you see the Installation Settings menu, press any key to halt the installation process.

Note: Carefully watch the boot sequence. The installation process does not provide much time to stop the installation process before automatically continuing.

3. Make sure that Installation is selected.
4. Enter the following kernel options in the boot options field:
`acpi=oldboot vga=normal`
5. Using the resolution function keys, select the screen resolution that matches your monitor.
6. Press **Enter** to continue the installation.
7. Read the SUSE End User License For SLES agreement and click **Accept**.
8. Select the language and click **Accept**.
9. When the type of installation pop-up window appears, select **New installation** and click **OK**.
10. Click **Change...**, and then click **Software**.
11. Change to **Default system for UnitedLinux**. If you are prompted to confirm that you really want to reset your detailed selection, click **Yes**.
12. Click **Detailed selection...**
13. From the selection pane, click **C/C++ compiler and tools** to add this package to the software list. Then click **Accept**.
14. Click **Partitioning** to create three new partitions.
 - a. Click **Create custom partition setup**, and then click **Next**.
 - b. Click the **Custom partitioning -- for experts**, and then click **Next**.
 - c. Select the **/dev/sda** disk and click **Delete** to remove all partitions. When prompted to confirm the deletion of all **/dev/sda** partitions, click **Yes**.

- d. Create the first partition:
 - 1) Click **Create**. If prompted about the disk on which to create the partition, select **/dev/sda**.
 - 2) When prompted about the type of partition to be created, select **primary partition**.
 - 3) Set up the partition using these values:
 - Starting cylinder: 0 (should already be set).
 - Ending cylinder: 1266
 - Format - filesystem: ReiserFS
 - Mount point: / (should already be set)
 - 4) Click **OK** to add the partition.
- e. Create the second partition:
 - 1) Click **Create**. If prompted about the disk on which to create the partition, select **/dev/sda**.
 - 2) When prompted about the type of partition to be created, select **primary partition**.
 - 3) Set up the partition using these values:
 - Starting cylinder: 1267
 - Ending cylinder: 1528
 - Format - filesystem: Swap
 - Mount point: swap
 - 4) Click **OK** to add the partition.
- f. Create the third partition:
 - 1) Click **Create**. If prompted about the disk on which to create the partition, select **/dev/sda**.
 - 2) When prompted about the type of partition to be created, select **primary partition**.
 - 3) Set up the partition using these values:
 - Starting cylinder: 1529
 - Ending cylinder: 3617
 - Format - filesystem: ReiserFS
 - Mount point: /var
 - 4) Click **OK** to add the partition.

The following table summarizes the partition settings.

Partition	Size	Filesystem	Mount	Cylinder Numbers
/dev/sda1	9.7 GB	ReiserFS	/	0 - 1266
/dev/sda2	2 GB	Swap	swap	1267 - 1528
/dev/sda3	16.0 GB	ReiserFS	/var	1529 - 3617

- g. Click **Next** to continue.
15. Click **Change...** to modify other settings, such as the timezone. After modifying these settings click **Accept**.
16. Click **Accept** to continue the installation.
17. At the warning prompt, click **Yes, install** to continue.
18. When prompted, insert the UnitedLinux CD-ROM 1 and click **OK**.
19. When prompted, insert UnitedLinux CD-ROM 2 in the CD-ROM drive and click **OK**.
20. When prompted, insert the SUSE SLES Version 8 CD-ROM 1 and click **OK**.

21. When prompted that the base system is installed, removed the CD-ROM and click **OK**.
22. The root user is created. When prompted, enter a password for root and click **Next**.
23. You are prompted to create more users. Create at least one additional user to prevent errors during the installation. For example, consider creating a user called **guest** if you do not need additional users. Click **Next** after you have created additional users.
24. You are prompted to either enable or disable the 3D graphics-capable card. Click **No** to disable the 3D graphics-capable card.
25. You are prompted to choose either graphical or text mode as the default desktop setting. Click **Text mode only** and click **Accept**.

Note: If you choose graphical mode as the default desktop setting, you disable the X Window System during the installation process. Text mode is the recommended setting.

26. When prompted about detecting printers, select **Skip detection**.
27. Click Network Interfaces to configure the interfaces.

Note: When you install SUSE LINUX Enterprise Server 8.0, it configures the primary interface as the first device it discovers in the following order (as you are looking at the engine from the rear):

- a. Top PCI adapter.
 - b. Bottom PCI adapter.
 - c. Left Ethernet port.
 - d. Right Ethernet port.
28. Verify that the primary interface (ETH0) is configured correctly by viewing the interface listed in the Already Configured pane.
 - a. If the incorrect interface is configured as the primary interface:
 - 1) Click **Change...**
 - 2) From the Network Cards Overview pane, make sure that the interface is selected.
 - 3) Click **Delete**.
 - 4) Click **Finish**.
 - b. To configure the correct interface.
 - 1) Select the primary interface from the list of interfaces in the Available Network Cards pane.
 - 2) Click **Configure**. For example, the first IBM 82546EB Gigabit Ethernet Controller is the left Ethernet port.
 - 3) Make sure the correct interface (network device) is selected.
 - 4) Click **Static IP address**.
 - 5) Fill in the IP address and subnet mask for this metadata server.
 - 6) Click **host name and name server**.
 - 7) Fill in the host name of the metadata server and the domain name.
 - 8) Fill in any name servers and domain search names based on your network configuration.
 - 9) Click **Next** twice. The configured network interface is displayed in the Already Configured Devices pane.
 - 10) Continue configure other interfaces as needed.
 - 11) Click **Finish**.

29. Click **Next** to complete the installation.

Disabling the automatic starting of the X Window System

This topic describes how to ensure that the X Window System does not start automatically when you boot the engine.

If you did not set text mode as the default desktop setting when you installed SUSE Linux Enterprise Server 8.0, you must disable it from automatically starting when you boot the engine.

1. Make sure that you are logged in as root.
2. Change to the `/etc` directory.
3. Edit `inittab`.
4. On the line, `id:5:initdefault`, change the 5 to a 3. The result should look like this:

```
id:3:initdefault
```
5. Save the file.
6. Reboot the engine.

Setting the time and date on the Metadata servers

This topic describes how to set the date and time on a metadata server.

Note: For the proper operation of the SAN File System, you must make sure that the system time and the hardware clock are synchronized on each metadata server engine in the cluster.

1. Log in as root.
2. Set the clock. For example:

```
# hwclock --set --date "Friday Sep 12 10:00"
```

Note: The time is set using a 24-hour format.

3. Set the time zone if you did not set it during the installation of the operating system. For example:

```
# rm /etc/localtime  
# ln -s /usr/share/zoneinfo/EST5EDT /etc/localtime
```
4. Set the system time from the hardware clock. For example:

```
# hwclock --hctosys
```

Apply United Linux Service Pack 3 (SP3) updates

This topic describes the procedure for applying the United Linux Service Pack 3 updates to a metadata server.

Apply the updates only after performing an initial installation of SUSE LINUX Enterprise Server 8.0 on the metadata server engine.

Tip: The process of applying the service pack updates might take a while to complete.

1. Insert the United Linux Service Pack CD-ROM into the CD-ROM drive.
2. Mount the CD-ROM:

```
mount /media/cdrom/
```
3. Run the installation script:

```
/media/cdrom/install.sh
```
4. Select **Option 1 - Update System to Service Pack 3 level**.
5. After the updates have been applied, you are prompted to quit. Press **Enter**.

6. Unmount the CD-ROM drive.
`umount /media/cdrom/`
7. Remove the CD-ROM from the drive.
8. Reboot the engine and log in as root.
9. Verify that the required kernel level is installed:
`rpm -qa | grep -e k_smp -e kernel`

The correct kernel level should be listed:

```
# rpm -qa |grep -e k_smp -e kernel
k_smp-2.4.21-138
kernel-source-2.4.21-138
```

Installing prerequisite software

This topic describes the software that must be loaded on the metadata server engine before installing the SAN File System software.

Note: You must be logged in as root to install software.

The following prerequisite software must be installed on the metadata server engine:

- QLogic driver. See “Install QLogic driver.”
- Management Processor Command Line Interface (MPCLI). See “Install MPCLI” on page 38.
- IBM Java™ Runtime Environment. See “Install the Java Runtime Environment” on page 38
- IBM Director Agent. See “Install the IBM Director agent” on page 38.
- Eclipse. See “Install Eclipse” on page 38.
- The `ibmusbasm` daemon. See “Install `ibmusbasm`” on page 38.
- `Openslp`. See “Install OpenSLP” on page 39.
- IBM Subsystem Device Driver (SDD). See “Install the IBM Subsystem Device Driver (SDD)” on page 39.
- IBM Websphere 5.0 Express. See “Install IBM WebSphere 5.0 Express” on page 39.

With the exception of the QLogic driver, all of this software is available on the SAN File System CD-ROM.

Install QLogic driver:

This topic describes the procedure for installing the QLogic driver.

Before you begin: Ensure that you have downloaded the QLogic driver `.tgz` file from the QLogic website and that you can use this driver with your storage subsystem.

1. Load the QLogic driver package into the temporary directory.
2. Unpack the QLogic `.tgz` file.
`cd /tmp`
`tar -xvzf qla2x00-v6.06.64-dist.tgz`
3. Run the `drvinstall` command:
`cd qlogic`
`./drvinstall`
4. Prepare the Linux kernel header files.

```

cd /usr/src/linux
make mrproper
cp /boot/config-2.4.21-138-smp .config
make oldconfig
make dep

```

5. Build the new driver and copy it to the appropriate directory:

```

cd /tmp/qlogic
make clean
make all SMP=1 OSVER=linux
make install

```

6. Make the QLogic driver load automatically when the engine is booted.
 - a. Edit /etc/sysconfig/kernel.

Tip: Make a backup of this file before you edit it.

- b. Add the QLogic driver to the INITRD_MODULES line. It should now look like this:

```
INITRD_MODULES="scsi_mod sd_mod mptscsih reiserfs qla2300"
```

- c. Save the file.
 - d. Run mkinitrd.

```
# mkinitrd
```

7. Reboot the engine.
8. Verify that the QLogic driver was successfully installed.

```
cat /proc/scsi/scsi
```

Result: All metadata volumes should now be visible. For example:

```

Attached devices:
Host: scsi0 Channel: 00 Id: 00 Lun: 00
  Vendor: LSILOGIC Model: 1030 IM          Rev: 1000
  Type: Direct-Access                    ANSI SCSI revision: 02
Host: scsi0 Channel: 00 Id: 08 Lun: 00
  Vendor: IBM      Model: 32P0032a S320 1 Rev: 1
  Type: Processor                      ANSI SCSI revision: 02
Host: scsi2 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 05
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 06
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 00 Lun: 07
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 2145           Rev: 0000
  Type: Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 01

```


Host: scsi3 Channel: 00 Id: 01 Lun: 07
Vendor: IBM Model: 2145 Rev: 0000
Type: Direct-Access ANSI SCSI revision: 03

Install MPCLI:

This topic describes how to install the Management Processor Command Line Interface (MPCLI).

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the MPCLI package:

```
rpm -Uvh /media/cdrom/mpcli-2.0-1.0.i386.rpm
```

Install the Java Runtime Environment:

This topic describes how to install the Java Runtime Environment 1.3.1-6.

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the IBM Java Runtime Environment package:

```
rpm -Uvh /media/cdrom/IBMJava2-JRE-1[1].3.1-6.0.i386.rpm
```

Depending on the shell that you are using, you may need to include a backslash (\) in front of the open and close bracket.

Note: If you receive a warning about a version of the package already being installed, you can ignore it. If the existing version was supplied by SUSE, it will have been packaged so that it installs under a different directory tree. To avoid unexpected results, use YaST2 to remove the SUSE-supplied version.

Install the IBM Director agent:

This topic describes how to install the IBM Director agent.

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Run the IBM Director Agent script:

```
/media/cdrom/IBMDirectorAgent4.11-1.sh
```

Note: If you receive the warning Starting Pegasus CIMOM ftp: connect: Network is unreachable, you can ignore it. Either the FRU is not supported or this system does not have access to the FRU file. Run the help command for getfru for more information (man getfru).

Install Eclipse:

This topic describes how to install Eclipse.

1. Make sure the SAN File System CD-ROM into the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the Eclipse package:

```
rpm -Uvh /media/cdrom/eclipse-2.0.2-1.i386.rpm
```

Install ibmusbasm:

This topic describes how to install ibmusbasm.

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the `ibmusbasm` package:


```
rpm -Uvh /media/cdrom/ibmusbasm-1.09-2.i386.rpm
```

A message similar to the following will be displayed to verify the installation.

```
Found Product ID 4001 USB Service Processor. Installing
the USB Service Processor Driver.
```

Install OpenSLP:

This topic describes how to install OpenSLP.

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the `openslp` package:


```
rpm -Uvh /media/cdrom/openslp-1.0.11-1.i386.rpm
```

Install the IBM Subsystem Device Driver (SDD):

This topic describes how to install the IBM Subsystem Device Driver (SDD).

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the SDD package:


```
rpm -Uvh /media/cdrom/IBMsdd-1.5.1.1-2.i686.u11.rpm
```
3. Start SDD.


```
sdd start
```

Install IBM WebSphere 5.0 Express:

This topic describes how to install IBM WebSphere[®] 5.0 Express.

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the WebSphere Express package:


```
rpm -Uvh /media/cdrom/bobcat-5.0.0-2.i386.rpm
```

If no errors occurred during installation, you will be returned to the shell prompt.

3. The default shell prompt displays the host name. If you are not using the default prompt or the host name is not displayed as part of the prompt, verify that the `HOSTNAME` variable is set.

```
# hostname
mds1
```

If the `HOSTNAME` variable is not set, perform these steps to set it.

- a. Edit the contents of `/etc/HOSTNAME`
 - b. Run this command:


```
export HOSTNAME=`cat /etc/HOSTNAME`
```
4. Change to the WebSphere Express source directory and install WebSphere 5.0 Express.


```
cd /opt/bobcat_src
./bobcat_install
```

Installing SAN File System software

This topic describes how to install the SAN File System software on a metadata server engine.

1. Make sure the SAN File System CD-ROM is in the CD-ROM drive and the CD-ROM drive is mounted.
2. Install the SAN File System package repository. This repository contains the software packages for all SAN File System components, including the metadata server, the administrative server, and all clients.

```
rpm -ivh /media/cdrom/sfs-package-build_level.i386.rpm
```

3. Change to the packages subdirectory.

```
cd /usr/tank/packages
```

4. Install the administrative package.

```
rpm -ivh sfs.admin.linux-build_level.i386.rpm
```

5. Install the metadata server package:

```
rpm -ivh sfs.server.linux-build_level.i386.rpm
```

6. Run the Target Machine Validation Tool (TMVT) to verify that your hardware and software prerequisites have been met.

```
/usr/tank/server/bin/tmvt -r report_file_name
```

Examine the results in *report_file_name*, paying particular attention to areas flagged as non-compliant. Resolve those prerequisites, and then rerun the tool until all prerequisites are compliant.

Note: TMVT non-compliance does not strictly prevent the installation of the SAN File System. It identifies deviations from the recommended hardware and software platform.

After you have completed the installation of the SAN File System software, you need to set your PATH environment variable to include the following paths:

- /usr/tank/server/bin
- /usr/tank/admin/bin

Setting up the master metadata server

This topic provides an overview of setting up the master metadata server.

1. Configure the RSA II adapter. See “Configuring the RSA II adapter.”
2. Upgrade the RSA II firmware, if needed. See “Upgrading RSA II firmware” on page 42
3. Run **setupsfs** to configure the master metadata server. See “Running the SAN File System setup utility” on page 42.

Configuring the RSA II adapter

This topic describes how to configure the RSA II adapter for the SAN File System.

Use the documentation supplied with the RSA II adapter card for configuration information. You need to change the factory-supplied IP address so that it is unique for your network. In addition, the SAN File System uses certain configuration settings (you will need to supply the configuration setting when you run the **setupsfs** command).

Note: You can use the **setupsfs** command to configure the RSA II adapter if it is not already configured. However, you must use this procedure to set the IP address, to set the user ID password to NULL, or to reset the user ID password.

This procedure assumes that you have properly cabled the RSA II adapter.

The RSA II adapters in the engines all come with the same default IP address. This address is: 192.168.70.125.

Note: If you have multiple RSA II adapters on the network at the same time, all will have the same IP address. If you use the Web interface to configure the RSA II adapter, you will not be able to determine the RSA II adapter that you are updating. Instead, you can update the IP address through the BIOS, and then use the Web interface to complete the configuration.

The recommended IP addresses for the RSA II adapters in the cluster are 192.168.70.1 through 192.168.70.*n*, where *n* is incremented by 1 for each additional RSA adapter. A cluster of eight engines would contain RSA adapters numbered from 192.168.70.1 to 192.168.70.8.

Number the engines and RSA adapters from 1 through *n*, starting with the top engine in the rack and ending with the bottom engine in the rack.

1. Reboot the engine.
2. Press **F1** when prompted to enter the BIOS setup panel.
3. Click **Advanced Setup**, and update the RSA II IP address and subnet mask.
4. Exit setup, saving your configuration changes.
5. After the system has finished booting, open a Web browser and point it to the IP address for the RSA II card.
6. Log on to the RSA II using the default user ID (USERID) and password (PASSWORD - the 0 is a zero).
7. In the left frame, click **Server --> ASM Control --> System Settings**.
8. Fill in the following information.
 - Name. The unique name of the RSA II adapter. This name must match the name of the metadata server. You will enter this name during metadata server setup.
 - IP address. Update the IP address to a unique IP address.
 - Server timeouts. Set the following timeouts:
 - Post watchdog. Set to 10 minutes.
 - OS watchdog. Set to 4 minutes.
 - Loader watchdog. Set to 10 minutes.
 - Set the date, time, and timezone.
 - Click **Save** to save your settings.
9. In the left frame, click **Server -->ASM Control -->Login Profiles**
10. Click on a "not used" link in the Login ID column.
11. Create a user ID and password to be used for logging in to the RSA II card. This user ID must have read/write (Supervisor) authority. You will enter this information during metadata server setup, and it is the same for all metadata server engines in the cluster.

Important: The password must contain only alphanumeric characters and it must be at least 5 characters long.

12. Click **Restart ASM** on the left frame.

13. You can verify the IP setting and new user ID by closing your browser, reopening it, and pointing to the new IP address.

Upgrading RSA II firmware

This topic describes how to upgrade the RSA II firmware for the metadata server engine.

You need to verify that your RSA II firmware is at version 1.06 (GEE834A). You can use the Web interface to the RSA II adapter to determine the firmware that you are currently using.

If not, you need to upgrade the RSA II firmware.

<http://www-307.ibm.com/pc/support/site.wss/document.do?lnodocid=MIGR-46489>

Make sure that you follow the instructions in the README to upgrade the RSA II firmware.

Running the SAN File System setup utility

This task describes how to use the setupsfs utility to set up the master metadata server.

All of the software should be loaded on the metadata server before you set it up as the master metadata server. The LDAP server should also be available. In addition, if you are using secured LDAP, the LDAP public certificate file should be copied to /usr/tank/admin.

The setupsfs utility is used to start the SAN File System metadata server configuration process.

1. Make sure that you are logged in as root.
2. Run the SAN File System setup utility.
 - a. If you are using system disks that have not previously been used with the SAN File System, run the setup utility as follows:

```
/usr/tank/admin/bin/setupfs -setmaster
```
 - b. If you are using system disks that were previously used with SAN File System and you want to completely re-initialize those disks, run the setup utility as follows:

```
/usr/tank/admin/bin/setupfs -setmaster -overwrite
```

Attention: You can use the -overwrite parameter to initialize the given master metadata server and system disks, regardless of whether they already contain cluster information. For example, if you get a failure in log.std indicating that the metadata disk is already labeled, but you are sure you wish to reuse that disk, you can rerun the setupsfs command with the -overwrite parameter. Remember that this parameter will destroy all metadata stored by the SAN File System.

Tip: If you are using ActiveDirectory as your LDAP server, you need to run setupsfs using the -debug parameter. Additional prompts will be displayed. Enter the appropriate information for the additional LDAP prompts, and accept the defaults for the additional prompts.

You will be prompted to enter the following information.

Table 1. Setupsfs prompts

Value	Description
SAN File System Server name	A unique name to be used for this metadata server engine. This name must be the same as the unique name used to configure the RSA II adapter on each engine.
SAN File System Cluster name	The name of the SAN File System cluster.
Server IP address	The IP address of the metadata server engine in dotted-decimal format.
Language	The language locale. For example, en_US.utf8
LDAP server IP address	The IP address of the LDAP server in dotted-decimal format.
LDAP user	The distinguished name of an authorized LDAP user. For example, cn=root
LDAP user password	The password for the LDAP user. This password must match the password set for this user in the LDAP server database.
LDAP secured connection	Set this value to true if you are using secured LDAP. Otherwise, set this value to false.
LDAP base distinguished name	The base distinguished name used to search for roles. For example, ou=NewRoles,o=ibm,c=US
LDAP member attribute	The attribute that contains the role members. For example, roleOccupant
LDAP certificate	If you are using secured LDAP, provide the fully qualified name of the LDAP certificate, which was obtained from the LDAP server. For example, /tmp/ldap.cert. Note: When you run setupsfs, this certificate will be embedded in the truststore.
RSA user name	The user ID used to access the RSA II adapter. The default is USERID.
RSA password	The password for the user ID. The default is PASSWORD (0 is zero).
CLI user	A user ID that has access to the administrative command-line interface. This user ID must be defined in the LDAP server database and must be set to the role of administrator.
CLI password	The password defined for the CLI user.
Truststore password	The password used to access the truststore.
Subordinate node list	A space-separated list of the IP addresses for subordinate metadata server engines in the cluster. For example, 192.168.10.69 192.168.10.79 192.168.10.89
Metadata disks	A space-separated list of raw devices on which SAN File System metadata is stored. For example, /dev/rvpathao /dev/rvpathap /dev/rvpathaq /dev/rvpathar. You can run the lsvpcfg command to decide which vpaths can be used for metadata disks. Then, you can specify the devices for this parameter using the following syntax: /dev/rvpath.

- When prompted to save the configuration, press **Enter**. You will copy this configuration to the subordinate metadata server engines when you set them up.

Installing subordinate metadata servers

This topic provides an overview of the steps required to install the SAN File System subordinate metadata servers.

Make sure that you have fulfilled the following prerequisites before installing each subordinate metadata server:

- You have unpacked the engine and installed the RSA II adapter card
- You have installed the engine in a rack.
- You have attached a keyboard, monitor, and mouse to the engine. Alternatively, a KVM as the console for the engine.

Use the documentation that comes with your engine and with the RSA II adapter to meet these prerequisites.

The procedures for installing the subordinate metadata servers are the same as the procedure for installing the master metadata servers. You need to perform these procedures on each subordinate metadata server that will be in the SAN File System cluster.

1. Prepare the engine for installation. This includes ensuring that the engine is properly cabled, making sure that you obtain all prerequisite software, and running the SAN File System mirroring utility. See “Preparing the subordinate engine for installation.”
2. Install all software, including the operating system, prerequisite software, and the SAN File System software. See “Installing software for the subordinate metadata server.”
3. Set up the subordinate metadata server by copying configuration files from the master metadata server, configuring the RSA II adapter, and running the SAN File System setup utility. See “Setting up a subordinate metadata server” on page 45.

Preparing the subordinate engine for installation

This topic describes how to prepare an engine to be installed.

The procedures to prepare a subordinate engine for installation are the same as the procedures you use to prepare a master engine for installation.

1. Ensure that the engine, including the RSA II adapter is cabled properly. See “Cabling” on page 26.
2. Obtain software to be used on the upgrade process. See “Obtain prerequisite software” on page 30.
3. Update the system BIOS, if necessary.
4. Use the LSI Logic Configuration Program documentation provided with the engine to mirror the boot drive.

Installing software for the subordinate metadata server

This topic provides an overview of the steps required to install the software for a subordinate metadata server.

The procedures to install software on the subordinate metadata server are the same as the procedures you use to install software on the master metadata server.

1. Install SUSE LINUX Enterprise Server 8.0. See “Installing the operating system” on page 31.

2. Disable the automatic starting of X-Windows (if you choose to have graphical mode as the default desktop setting during the installation of the operating system). See “Disabling the automatic starting of the X Window System” on page 34.
3. Set the time and date on the metadata server engine. See “Setting the time and date on the Metadata servers” on page 34.
4. Apply the United Linux Service Pack 3. See “Apply United Linux Service Pack 3 (SP3) updates” on page 34.
5. Install all prerequisite software, such as the QLogic driver and the IBM Director agent. See “Installing prerequisite software” on page 35.
6. Install the SAN File System software. See “Installing SAN File System software” on page 40.

Setting up a subordinate metadata server

This topic provides an overview of setting up a subordinate metadata server.

The procedures for setting up a subordinate metadata server are the same as the procedures you use to set up the master metadata server, with one exception. When you run the setup utility, you use slightly different parameters.

You need to perform these procedures on all subordinate metadata server engines.

1. Copy `tank.properties` and the truststore. See “Copying `tank.properties` and the truststore.”
2. Configure the RSA II adapter. See “Configuring the RSA II adapter” on page 40.
3. Upgrade the RSA II firmware, if needed. See “Upgrading RSA II firmware” on page 42
4. Run the `setupfs` utility to configure the subordinate metadata server. See “Running the SAN File System setup utility.”

Copying `tank.properties` and the truststore

This topic describes how to copy `tank.properties` and the truststore from the master metadata server to a subordinate metadata server.

1. From the master metadata server, copy `tank.properties` to the subordinate metadata server.

```
scp /usr/tank/admin/config/tank.properties
userID@metadata_server_name:/usr/tank/admin/config/tank.properties
```

For example:

```
scp /usr/tank/admin/config/tank.properties
root@sub_mds1:/usr/tank/admin/config/tank.properties
```

Attention: When you copy `tank.properties` and then run `setupfs` on the subordinate metadata server, the IP address and server name of the master metadata server will be displayed as a default. Make sure that you change these values to match the appropriate values for the subordinate metadata server.

2. Copy the truststore file to the subordinate metadata server. If you are using secured LDAP, the truststore file will also contain the LDAP certificate.

```
scp /usr/tank/admin/truststore
userID@metadata_server_name:/usr/tank/admin/truststore
```

Running the SAN File System setup utility

This topic describes how to use the `setupfs` utility to set up subordinate metadata servers.

You will need to perform these steps on all subordinate metadata servers in the cluster.

1. Make sure that you are logged in as root.
2. Run the SAN File System setup utility.
`/usr/tank/admin/bin/setupfs`

When prompted, enter the new name and IP address for this metadata server. Accept the defaults for all other prompts.

3. Save the configuration by pressing **Enter**.

Setting up the cluster

This topic provides an overview of the steps required to set up the SAN File System cluster.

1. Start the cluster. See “Forming the cluster.”
2. Validate that the cluster has been installed successfully. See “Validating cluster installation.”

Forming the cluster

This topic describes how to form the new cluster.

Before forming the cluster, you should have set up the master metadata server and all subordinate metadata servers.

1. Log in to the master metadata server as root.
2. Form the new cluster.

```
# /usr/tank/admin/bin/setupfs -newcluster
```

Validating cluster installation

This topic describes how to validate that the cluster was installed correctly.

1. Make sure that you are logged in to the master metadata server as root.
2. List all servers in the SAN File system cluster.

```
/usr/tank/admin/bin/sfsccli lserver
```

You should see all metadata servers in the list.

Name	State	Server Role	Filesets	Last Boot
mstr-mds	Online	Master		3 Feb 13, 2004 2:46:28 PM
sub1-mds	Online	Subordinate		1 Feb 16, 2004 4:26:08 AM

Setting up clients

This topic describes the general process for setting up clients.

You can set up SAN File System clients running on the following operating system:

- Windows 2000 Server and Advanced Server. See “Installing SAN File System on a Windows client” on page 47.
- IBM AIX Version 5.1 (32-bit) and IBM AIX Version 5.2 (32-bit and 64-bit). See “Installing SAN File System on an AIX client” on page 49.
- Linux Red Hat Advanced Server 3.0. See “Installing SAN File System on a Linux client” on page 51.

- Sun Solaris 9 (64-bit). See “Installing SAN File System on a Solaris client” on page 53.

Optionally, you can install IBM Subsystem Device Driver (SDD) on the SAN File System clients for multipathing support. See “Installing SDD on clients.”

Installing SDD on clients

This topic explains where to go for information about installing the IBM Subsystem Device Driver (SDD) on SAN File System clients.

The IBM Subsystem Device Driver (SDD) provides the multipath configuration environment support for a host system that is using an IBM TotalStorage SAN File System. SDD is optional on SAN File System clients.

Note: If you install SDD on SAN File System clients, you should install the latest SDD (currently version 1.5.1). In addition, it must be installed before you install the SAN File System software on the client.

Refer to the SDD User’s Guide (SC26-7637) for installation procedures. You can find this document at www.ibm.com/storage/support. Choose **Subsystem Device Driver** under the **Storage Software** option. Then click **Documentation** under the Information heading.

Installing SAN File System on a Windows client

This topic provides the general procedure for installing the SAN File System on a Windows client. These steps must be performed on each Windows client in the SAN File System.

1. Obtain the SAN File System client software. See “Obtain version 2.1 software for a Windows client.”
2. Prepare the Windows client for installation by stopping all applications on the client. Refer to the documentation that comes with the application for information about stopping it.
3. Install the client software. See “Installing the SAN File System software” on page 48.
4. Validate the installation. See “Validating the installation of SAN File System on a Windows client” on page 48.

Obtain version 2.1 software for a Windows client

This topic explains how to obtain the version 2.1 SAN File System software for a Windows client.

The client installation package is called `sfs-client-WIN2K-version.exe`. You can load this package on the client from the SAN File System package repository. The package repository is located on each metadata server engine. Use the SAN File System console to transfer the executable from a metadata server engine.

To transfer the executable using the SAN File System console:

1. Start the SAN File System console from the client.
2. Select **Download Client Software**.
3. Follow the prompts to save the executable to a temporary directory.

Installing the SAN File System software

This topic describes how to install version 2.1 of the SAN File System on a Windows client.

- The client for Windows can be installed only on Windows 2000 Server and Advanced Server. A minimum of Service Pack 4 is required. The operating system must already be installed with the appropriate service packs.
 - The client for Windows requires least 10 MB of free disk space.
 - You must have Administrator privileges to install the client for Windows.
 - A SAN File System client can be attached to one SAN File System server cluster only.
 - The LAN and SAN should be installed and configured as well as prerequisite products such as IPSec, FC-HBA drivers, networking, fibre-channel switch firmware, and storage device firmware.
 - There must be a free drive letter.
 - If you install the client during a Windows Terminal Service (WTS) session, the drive letter assigned to the file system will be visible only to that WTS session (private name space). To globally share the file system on WTS, reboot the client system.
 - The metadata server must be up and running with the IP address and port defined. This information is needed during setup.
 - Some basic startup, shutdown, or error messages are written to the Windows system log (event viewer).
1. Navigate to the directory where the windows client software is located.
 2. Run the setup command.
`sfs-client-WIN2K-build_level.exe`
 3. Select the language that you want to use for the installation process, and click **OK**.
 4. The Welcome window appears. Click **Next**.
 5. The SAN File System client settings panel is displayed. Fill in the configuration information.
 - SAN File System server name (no default)
 - SAN File System server port (default is 1700)
 - SAN File System preferred drive letter (default is T:)
 - SAN File System client name (default is the short version of the hostname)
 - SAN File System network connection type (default is TCP)
 - SAN File System client critical error handling policy (default is log)

Note: Make sure that you check the box **Disable Disk Management Write Signature**.

6. Click **Next** to continue.
7. The Start Copying Files panel is displayed. Verify that the settings are correct and click **Next**.
8. When prompted to start the SAN File System client now, click **Yes**.
9. Click **Finish**.

Validating the installation of SAN File System on a Windows client

This topic describes how to validate that version 2.1 of the SAN File System was installed properly on a Windows client.

1. Open Windows Explorer and verify that the drive letter you specified in the configuration is listed.

Note: If the SAN File System client is not started, you can start it from a command-prompt window:

```
net start stfs
```

2. If you do not see the drive letter listed, reboot the client.

Installing SAN File System on an AIX client

This topic provides the general steps for installing the SAN File System on an AIX client. These steps must be performed on each AIX client in the SAN File System.

1. Obtain the SAN File System client software. See “Obtain version 2.1 software for an AIX client.”
2. Prepare the AIX client for upgrading by stopping all applications on the client. Refer to the documentation that comes with the application for information about stopping it.
3. Install the client software. See “Installing the SAN File System software” on page 50.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on an AIX client” on page 51.

After you have completed the installation of the AIX client, you need to set your PATH environment variable to include the following paths:

- /usr/tank/client/bin
- /usr/tank/migration/bin

Obtain version 2.1 software for an AIX client

This topic explains how to obtain the version 2.1 SAN File System software for an AIX client.

If you are installing the client on AIX version 5.1, the package name is sfs.client.AIX51. If you are installing the client on AIX version 5.2, the package name is sfs.client.AIX52.

You can load either of these packages on the client from the SAN File System package repository. The package repository is located on each metadata server engine. Use either ftp or the SAN File System console to transfer the package from a metadata server engine.

To transfer the package using ftp:

1. On the master metadata server, change directories to the package repository.

```
cd /usr/tank/packages
```
2. Set up an ftp session with the client.

```
ftp client_IP_address
```
3. Log in to the client.
4. Set the transfer mode to binary

```
bin
```
5. Change directories to a temporary directory on the client

```
cd /tmp
```
6. Transfer the package to the client. For example, to transfer the SAN File System client software for AIX 5.2:

```
mput sfs.client.AIX52
```
7. When prompted to confirm that you want to put the file, type Y.
8. Exit ftp

bye

To transfer the package using the SAN File System console:

1. Start the SAN File System console from the client.
2. Select **Download Client Software**.
3. Follow the prompts to save the package to a temporary directory.

Installing the SAN File System software

This topic describes how to install version 2.1 of the SAN File System on an AIX client

- You can install the SAN File System client on the following AIX operating systems:
 - AIX 5.1 (32-bit only) uniprocessor or multiprocessor with maintenance level 3.

Note: For AIX 5.1, the bos.mp (multiprocessor) or bos.up (uniprocessor) packages must be at least 5.1.0.58 or higher. AIX 5.1 32-bit high availability cluster multi-processing (HACMP™) environments are supported at the specified maintenance level.

- AIX 5.2 (32-bit and 64-bit) with bos.mp 5.2.0.18.
 - You must have root privileges to install the client for AIX.
1. Enable asynchronous input/output on the AIX client, which is required for the SAN File System.
 - a. Log on to the client as root.
 - b. Start **smit**.
 - c. Select **Devices**
 - d. Select **Asynchronous I/O**
 - e. Select **Asynchronous I/O (Legacy)**

Note: You will see this choice only if you are disabling asynchronous I/O on AIX 5.2. If you are using AIX 5.1, skip this step.

- f. Select **Change/Show Characteristics of Asynchronous I/O**
 - g. Use the tab key to set the STATE to be configured at system restart to **available**.
 - h. Press Enter.
 - i. Exit smit.
 - j. Run `cfgmgr` to apply the changes.
2. Navigate to the directory where the client installation package is located.
 3. Optionally, use `installp` or `smit` to preview the installation of the package. This way, you can resolve any warnings before you actually commit the installation.

```
installp -ap -d . client_package_name
```
 4. Use `installp` or `smit` to install the package. For example:

```
inutoc .  
installp -ac -d . client_package_name
```
 5. Configure and start the client. Run the `setup` command with the `-prompt` parameter.

```
/usr/tank/client/bin/setupstclient -prompt
```

You will be prompted to enter values for the client configuration.

- SAN File System server name (no default)
- SAN File System server port (default is 1700)

- SAN File System mount point (no default)
- SAN File System client name (default is the short version of the hostname)
- SAN File System network connection type (default is TCP)
- SAN File System client critical error handling policy (default is log)

In most cases you can accept the defaults. If you change these values, make sure that you type the new values correctly.

Tip: If you have installed SDD on the client, you should use the following device pattern when prompted for storage devices.

```
pat=/dev/rvpath*
```

Validating the installation of SAN File System on an AIX client

This topic describes how to validate that the SAN File System was installed properly on an AIX client.

1. Use the cat command

```
cat /usr/tank/client/VERSION
```

The results should be similar to the following:

```
VERSION 2.1.0
RELEASE 19
INTERFACE 0
```

2. Use the mount command to verify that the SAN File System is mounted on the client. The mount point for the SAN File System should be displayed.

Installing SAN File System on a Linux client

This topic provides the general steps for installing the SAN File System on a Linux client. These steps must be performed on each Linux client in the SAN File System.

1. Obtain the SAN File System client software. See “Obtain version 2.1 software for a Linux client.”
2. Prepare the Linux client for installation by stopping all applications on the client. Refer to the documentation that comes with the application for information about stopping it.
3. Install the client software. See “Installing the SAN File System software” on page 52.
4. Validate the installation. See “Validating the installation of SAN File System on a Linux client” on page 52.

Obtain version 2.1 software for a Linux client

This topic explains how to obtain the version 2.1 SAN File System software for a Linux client.

The client installation package is called `sfs.client.linux_RHEL3_9-build_level.i386.rpm`. You can load this package on the client from the SAN File System package repository. The package repository is located on each metadata server engine. Use either `scp` or the SAN File System console to transfer the package from a metadata server engine.

To transfer the package using secured copy:

1. On the client, change directories to a temporary directory.

```
cd /tmp
```

2. Copy the software package from the master metadata server.

```
scp userID@server_host_name:/usr/tank/packages/clientpackage.rpm .
```

For example:

```
scp root@mstr_mds:/usr/tank/packages/sfs.client.linux_RHEL3_9-  
build_level.i386.rpm
```

If you are using a graphical interface on Linux, you can transfer the package using the SAN File System console:

1. Start the SAN File System console from the client.
2. Select **Download Client Software**.
3. Follow the prompts to save the package to a temporary directory.

Installing the SAN File System software

This topic describes how to install version 2.1 of the SAN File System on a Linux client

You can install the SAN File System client on the Red Hat Enterprise Linux Advanced Server 3.0, 2.4.21-9.ELhugemem config for i686 or 2.4.21-138-smp.

1. Navigate to the directory where the client installation package is located.
2. Install the client package.

```
rpm -ihv sfs.client.linux_RHEL3_9-build_level.i386.rpm
```

3. Make sure that the master metadata server is running.
4. Configure and start the client. Run the setup command.

```
/usr/tank/client/bin/setupstclient -prompt
```

You will be prompted to enter values for the client configuration.

- SAN File System server name (no default)
- SAN File System server port (default is 1700)
- SAN File System mount point (no default)
- SAN File System client name (default is the short version of the hostname)
- SAN File System network connection type (default is TCP)
- SAN File System client critical error handling policy (default is log)

In most cases you can accept the defaults.

Tip: If you have installed SDD on the client, you should use the following device pattern when prompted for storage devices.

```
pat=/dev/vpath*[a-z]
```

Validating the installation of SAN File System on a Linux client

This topic describes how to validate that the SAN File System was installed properly on an Linux client.

1. Use the cat command

```
cat /usr/tank/client/VERSION
```

The results should be similar to the following:

```
VERSION 2.1.0  
RELEASE 19  
INTERFACE 0
```

2. Use the mount command to verify that the SAN File System is mounted on the client. The mount point for the SAN File System should be displayed.

Installing SAN File System on a Solaris client

This topic provides the general steps for installing the SAN File System on a Solaris client. These steps must be performed on each Solaris client in the SAN File System.

1. Obtain the SAN File System client software. See “Obtain version 2.1 software for a Solaris client.”
2. Prepare the Solaris client for installation by stopping all applications on the client. Refer to the documentation that comes with the application for information about stopping it.
3. Install the client software. See “Installing the SAN File System software.”
4. Validate the installation. See “Validating the installation of SAN File System on a Solaris client” on page 54.

Obtain version 2.1 software for a Solaris client

This topic explains how to obtain the version 2.1 SAN File System software for a Solaris client.

The client installation package is called `sfs.client.solaris9.build_level`. You can load this package on the client from the SAN File System package repository. The package repository is located on each metadata server engine. Use either `scp` or the SAN File System console to transfer the package from a metadata server engine.

To transfer the package using secured copy:

1. On the client, change directories to a temporary directory.

```
cd /tmp
```

2. Copy the software package from the master metadata server.

```
scp userID@server_host_name:/usr/tank/packages/clientpackage .
```

For example:

```
scp root@mstr_mds:/usr/tank/packages/sfs.client.solaris9.build_level
```

If you are using a graphical interface on Solaris, you can transfer the package using the SAN File System console:

1. Start the SAN File System console from the client.
2. Select **Download Client Software**.
3. Follow the prompts to save the package to a temporary directory.

Installing the SAN File System software

This topic describes how to install version 2.1 of the SAN File System on a Solaris client

- You can install the SAN File System client on the Sun Solaris 9 operating system.
- To install the SAN File System on a Solaris client, you must be logged on with root privileges.

1. Navigate to the directory where the client installation package is located.
2. Install the client package.

```
pkgadd -d sfs.client.solaris9.build_level IBMSFS
```

3. Enter All (the default) when prompted to select the packages to be installed.
4. Enter y when prompted to continue the installation.
5. Make sure that the master metadata server is running.
6. Configure and start the client.
 - a. Run the setup command

```
/usr/tank/client/bin/setupstclient -prompt
```

You will be prompted to enter values for the client configuration.

- SAN File System server name (no default)
- SAN File System server port (default is 1700)
- SAN File System mount point (no default)
- SAN File System client name (default is the short version of the hostname)
- SAN File System network connection type (default is TCP)
- SAN File System client critical error handling policy (default is log)

In most cases you can accept the defaults.

Validating the installation of SAN File System on a Solaris client

This topic describes how to validate that the SAN File System was installed properly on a Solaris client

1. Use the cat command

```
cat /usr/tank/client/VERSION
```

The results should be similar to the following:

```
VERSION 2.1.0  
RELEASE 19  
INTERFACE 0
```

2. Use the mount command to verify that the SAN File System is mounted on the client. The mount point for the SAN File System should be displayed.

Configuring SAN File System

This topic provides an overview of configuring the SAN File System.

There are several configuration steps required before you can begin using the SAN File System. These steps include:

- Configuring each of the metadata server engines in the cluster to use the master console as an SNMP manager. This allows you to enable the Service Alert and Remote Access features. See “Configuring metadata servers for SNMP traps.”
- Configure storage pools for storing both user data (user storage pools) and metadata (system storage pool). See “Creating storage pools” on page 55.
- Configure filesets to specify which metadata server will manage a particular fileset and how much space can be allocated to files within a fileset. See “Configuring filesets” on page 55.
- Create placement policies to specify how files will be placed in storage pools. See “Placement policies” on page 59.
- Optionally, you can migrate existing data to be managed by the SAN File System. See “Migrating data” on page 65.

Configuring metadata servers for SNMP traps

This topic describes how to configure metadata servers for service alerts by setting up the master console as an SNMP manager.

1. Add the master console as an SNMP manager.

```
/usr/tank/admin/bin/sfsccli addsnmpmgr -ip master_console_IP_address
```

2. Set the types of events that will generate an SNMP trap.

```
/usr/tank/admin/bin/sfsccli settrap -event sev
```

Your SAN File System will now send service alerts for metadata server failures.

Creating storage pools

This topic describes how to create storage pools.

The following prerequisites must exist before you can configure storage pools:

- There must be volumes available to create a new storage pool. If not:
 - The SAN File System administrator must request volumes from a storage device administrator.
 - The administrator of the storage device must assign Metadata server engines and SAN File System clients as hosts. In addition, the administrator must allocate LUNs to be used as devices in the system storage pool to metadata server engines and LUNs to be used as data devices to the appropriate clients.
- Only one system storage pool can exist.
- Make sure that you add LUNs to the default storage pool.
- You must be an administrator or IBM support representative to perform this task.

Perform these steps to create storage pools.

1. Using your Web browser, connect to the SAN File System Console.
2. In the My Work frame, click **Manage Storage** → **Create a Storage Pool**. View the list of steps to create a storage pool and click **Next**.
3. Under **Pool Settings** in the Create a Storage Pool panel, fill in the following fields:
 - Name of the new storage pool. You can enter up to 256 characters for a name, but the name must not currently exist.
 - Description of the new storage pool. You can enter up to 256 characters for a description.
 - Optionally, fill in these fields:
 - Logical partition size.
 - Allocation Size.
 - Usage threshold.
4. Click **Next** to continue.
5. Under **Select Client**, select a client and a fetch method to gather the available LUNs information for the next step, adding volumes to the storage pool. The default fetch method is to gather the LUN information from cache; you could also choose to rediscover the LUNs by selecting the Rediscover button. Click **Next**.
6. Under **Add Volumes**, select the LUNs to be added from the table. Click **Next**.
7. Under **Volume Settings**, fill in the **Volume Name Prefix** field, and click **Next**.
8. Verify your settings, and then click **Finish**.

Configuring filesets

This topic describes how to configure filesets.

Fileset quotas provide a way for an administrator to specify how much space can be allocated to files within a specific fileset. The default value is set to allow unlimited capacity; however, you can specify an alert value for the fileset. If a quota value is specified, the default alert value is 80%. If no quota value is specified, the default alert value is 0 (no alerts). When the space allocated to files within the fileset reaches the percentage of the quota, as specified by the alert

value, an SNMP alert is sent to the administrator. The administrator can also specify whether to use a hard or soft quota. If a hard quota is specified, allocations that cause a quota violation will fail and an SNMP alert will be sent. If a soft quota is specified, then the allocation will be allowed to succeed and an SNMP alert will be issued.

Filesets are statically bound to a Metadata server. When filesets are created, the GUI or CLI specifies the Metadata server name to which to bind the fileset along with other parameters. You can change fileset binding only by using the GUI or CLI command `chfilesetserver`.

You can reassign a fileset using these methods only:

- The cluster mode to Administrative if the Metadata server serving the fileset is in online mode and fileset is attached.
- The Metadata server serving the fileset is out of group, in which case you must certify that original Metadata server is offline by switching off the engine to avoid rogue Metadata server issues before moving the fileset from that Metadata server.

The Metadata server workload is not dynamically balanced by moving the filesets around the Metadata servers. All the filesets assigned to the offline Metadata server are inaccessible until that Metadata server comes online or you reassign those filesets to the online Metadata server.

1. In the My Work pane, click **Manage Filing**.
2. Click **Create a Fileset**.
3. In the Create a Fileset pane, fill in the Name and Description fields, and choose a server.
4. Under Attach Point, fill in the fields for Existing Directory Path, and New Directory Path. Other fields on this page are optional.

Creating a fileset for AIX

Perform™ the following steps to create a fileset for AIX.

1. In the My Work frame, click **Manage Filing** → **Create a Fileset**.
 - In the Create a Fileset panel:
 - a. Fill in the **Name** field (AIX_Fileset), the **Description** field (for example, A fileset for AIX-only files), and select a server (for example, ST0) from the drop-down list.
 - b. Under **Attach Point**, fill in the **Directory Path** field (for example, sanfs) and the **Directory Name** field (for example, aix51). Click **OK**.
 - Optionally, select a **Server Assignment Method** and **Quota Options**.
2. Click **Manage Filing** → **Filesets**. Verify your new fileset in the list.
3. Grant root privileges to the client by clicking **Manage Servers and Clients** → **Client Sessions**.
 - a. In the Client Sessions panel, Select a client, select **Grant Clients Root Privileges** from the drop-down list, and then click **Go**.
4. On the IBM AIX client machine, switch to the SAN File System mount point, and change to the global fileset directory.

```
# pwd
/mnt/SAN_FS_MOUNTPT/sanfs
# ls
```

```
total 8
d----- 2 1000000 1000000 4096 July 3 10:21 aix51
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 .flashcopy
#|
```

5. Change the ownership and permissions of the fileset.

```
# chown root:system aix51
# chmod 755 aix51
# ls
total 8
drwxr-xr-x 2 root    system  4096 July 3 10:21 aix51
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 .flashcopy
#|
```

Creating a fileset for Linux

Perform the following steps to create a fileset for Red Hat Advanced Server 2.1 Linux.

1. In the My Work frame, click **Manage Filing** → **Create a Fileset**.
 - In the Create a Fileset panel:
 - a. Fill in the **Name** field (Linux_Fileset), the **Description** field (for example, A fileset for Linux-only files), and select a server (for example, ST0) from the drop-down list.
 - b. Under **Attach Point**, fill in the **Directory Path** field (for example, sanfs) and the **Directory Name** field (for example, linux21). Click **OK**.
 - Optionally, select a **Server Assignment Method** and **Quota Options**.
2. Click **Manage Filing** → **Filesets**. Verify your new fileset in the list.
3. Grant root privileges to the client by clicking **Manage Servers and Clients** → **Client Sessions**.
 - a. In the Client Sessions panel, Select a client, select **Grant Clients Root Privileges** from the drop-down list, and then click **Go**.
4. On the Red Hat Linux client machine, switch to the SAN File System mount point, and change to the global fileset directory.

```
# pwd
/mnt/SAN_FS_MOUNTPT/sanfs
# ls
total 8
d----- 2 1000000 1000000 4096 July 3 10:21 linux
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 .flashcopy
#|
```

5. Change the ownership and permissions of the fileset.

```
# chown root: linux
# chmod 755 linux
# ls
total 8
drwxr-xr-x 2 root    root    4096 July 3 10:21 linux
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 .flashcopy
#|
```

Creating a fileset for Solaris

Perform the following steps to create a fileset for Solaris.

1. In the My Work frame, click **Manage Filing** → **Create a Fileset**.
 - In the Create a Fileset panel:
 - a. Fill in the **Name** field (Solaris_Fileset), the **Description** field (for example, A fileset for Solaris-only files), and select a server (for example, ST0) from the drop-down list.

- b. Under **Attach Point**, fill in the **Directory Path** field (for example, sanfs) and the **Directory Name** field (for example, solaris9). Click **OK**.
 - Optionally, select a **Server Assignment Method** and **Quota Options**.
2. Click **Manage Filing** → **Filesets**. Verify your new fileset in the list.
3. Grant root privileges to the client by clicking **Manage Servers and Clients** → **Client Sessions**.
 - a. In the Client Sessions panel, Select a client, select **Grant Clients Root Privileges** from the drop-down list, and then click **Go**.
4. On the Solaris client machine, switch to the SAN File System mount point, and change to the global fileset directory.

```
# pwd
/mnt/SAN_FS_MOUNTPT/sanfs
# ls
total 8
d----- 2 1000000 1000000 4096 July 3 10:21 solaris
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 .flashcopy
#|
```

5. Change the ownership and permissions of the fileset.

```
# chown root: solaris
# chmod 755 solaris
# ls
total 8
drwxr-xr-x 2 root root 4096 July 3 10:21 solaris
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 .flashcopy
#|
```

Creating a fileset for Windows

Perform the following steps to create a fileset for Windows.

1. In the My Work frame, click **Manage Filing** → **Create a Fileset**.
 - In the Create a Fileset panel:
 - a. Fill in the **Name** field (for example, Win_Fileset), the **Description** field (for example, A fileset for Windows files), and select a server (for example, ST0) from the drop-down list.
 - Optionally, select a **Server Assignment Method** and **Quota Options**.
 - a. Under **Attach Point**, fill in the **Directory Path** field (for example, sanfs) and the **Directory Name** field (for example, win2k), and then press **OK**.
2. Click **Manage Filing** → **Filesets**. Verify your new fileset in the list.
3. Grant root privileges to the client by clicking **Manage Servers and Clients** → **Client Sessions**.
 - a. In the Client Sessions panel, select a client, select **Grant Clients Root Privileges** from the drop-down list, and then click **Go**.
4. Open Microsoft® Windows Explorer, expand the SAN File System drive letter, and then select the fileset you just created (for example, win2k). Set the owner by right-clicking and selecting **Properties**. Click the **Security** tab. Click **Advanced**, and then click the **Owner** tab.
5. Set permissions by selecting the folder containing the fileset (for example, win2k) and right-clicking and selecting **Properties**. Click the **Security** tab. Click **Advanced**, and then click the **Permissions** tab. Select a permission, and then click **Apply** and **OK**. Click **OK** again.

When you are done, the fileset is now ready for use on the Windows 2000 operating system.

SETSTGPOOL *'pool_name'*

Identifies the pool in which you want to place all files that match the rule criteria (fileset and SQL expression).

FOR FILESET (*'fileset_name'*)

Identifies one or more filesets in which the file is created to determine where the file is to be placed. In the case of nested filesets, the rules apply if the file is created in the innermost fileset.

where *SQL_expression*

Compares the file attributes specified in the rule with the attributes of the file being created to determine where the file is to be placed. The *SQL_expression* can be any combination of standard SQL-syntax expressions, including comparison predicates, between predicates, in predicates, like predicates, mathematical value expressions, and boolean, string and numeric literals.

Note: Case expressions and compared-when clauses are not allowed.

SAN File System supports built-in functions, which can be used in comparison predicates, between predicates, in predicates, and like predicates. These functions are organized in three categories: date and time manipulation, numeric calculations, and string manipulation.

Attributes

You can use any of these attributes in the expression:

NAME

Name of the file. You can use a % wildcard in this name to represent zero or more characters and use the _ wildcard to represent a single character.

CREATION_DATE

Date and time that the file was created.

GROUP_ID

Numeric group ID. This attribute is valid only for AIX clients.

USER_ID

Numeric user ID. This attribute is valid only for AIX clients.

String functions

You can use these string-manipulation functions on file names and literals.

Note: You must enclose strings in single-quotation marks. You may include a single-quotation mark in a string by using two single-quotation marks (for example, 'a''b' represents the string a'b).

CHAR(x)

Converts an integer *x* to a string.

CHARACTER_LENGTH(x)

Determines the number of characters in string *x*.

CHAR_LENGTH(x)

Determines the number of characters in string *x*.

CONCAT(x,y)

Concatenates strings *x* and *y*.

HEX(x)

Converts an integer *x* in hexadecimal format.

LCASE(*x*)

Converts string *x* to lowercase.

LEFT(*x,y,z*)

Left justifies string *x* in a field of *y* characters, optionally padding with character *z*.

LENGTH(*x*)

Determines the length of the data type of string *x*.

LOWER(*x*)

Converts string *x* to lowercase.

LTRIM(*x*)

Removes leading blank characters from string *x*.

POSITION(*x* IN *y*)

Determines the position of string *x* in string *y*.

POSSTR(*x,y*)

Determines the position of string *y* in string *x*.

RIGHT(*x,y,z*)

Right justifies string *x* in a field of *y* characters, optionally padding with character *z*.

RTRIM(*x*)

Removes the trailing blank characters from string *x*.

SUBSTR(*x* FROM *y* FOR *z*)

Extracts a portion of string *x*, starting at position *y*, optionally for *z* characters (otherwise to the end of the string).

SUBSTRING(*x* FROM *y* FOR *z*)

Extracts a portion of string *x*, starting at position *y*, optionally for *z* characters (otherwise to the end of the string).

TRIM(*x*)

Trims blank characters from the beginning and end of string *x*.

TRIM(*x* FROM *y*)

Trims blank characters that are *x* (LEADING, TRAILING, or BOTH) from string *z*.

TRIM(*x y* FROM *z*)

Trims character *y* that is *x* (LEADING, TRAILING, or BOTH) from string *z*.

UCASE(*x*)

Converts the string *x* to uppercase.

UPPER(*x*)

Converts the string *x* to uppercase.

Numerical functions

You can use these numeric-calculation functions to place files based on either numeric parts of the file name, numeric parts of the current date, and AIX-client user IDs or group IDs. These can be used in combination with comparison predicates and mathematical infix operators (such as addition, subtraction, multiplication, division, modulo division, and exponentiation).

INT(*x*)

Converts number *x* to a whole number, rounding up fractions of .5 or greater.

INTEGER(*x*)

Converts number *x* to a whole number, rounding up fractions of .5 or greater.

MOD(*x,y*)

Determines $x \% y$.

Date and time functions

You can use these date-manipulation and time-manipulation functions to place files based on when the files are created at the client and the local time of the subordinate metadata server serving the directory within which the file is being created.

CURRENT DATE

Determines the current date on the subordinate metadata server.

CURRENT_DATE

Determines the current date on the subordinate metadata server

CURRENT TIME

Determines the current time on the subordinate metadata server.

CURRENT_TIME

Determines the current time on the subordinate metadata server.

CURRENT_TIMESTAMP

Determines the current date and time on the subordinate metadata server.

CURRENT_TIMESTAMP

Determines the current date and time on the subordinate metadata server.

DATE(*x*)

Creates a date out of *x*.

DAY(*x*)

Creates a day of the month out of *x*.

DAYOFWEEK(*x*)

Creates the day of the week out of date *x*, where *x* is a number from 1 to 7 (Sunday=1).

DAYOFYEAR(*x*)

Creates the day of the year out of date *x*, where *x* is a number from 1 to 366.

DAYS(*x*)

Determines the number of days since 0000-00-00.

DAYSINMONTH(*x*)

Determines the number of days in the month from date *x*.

DAYSINYEAR(*x*)

Determines the day of the year from date *x*.

HOUR(*x*)

Determines the hour of the day (a value from 0 to 23) of time or timestamp *x*.

MINUTE(*x*)

Determines the minute from date *x*.

MONTH(*x*)

Determines the month of the year from date *x*.

QUARTER(*x*)

Determines the quarter of year from date *x*, where *x* is a number from 1 to 4 (for example, January, February, and March is quarter 1).

SECOND(*x*)

Returns the seconds portion of time *x*.

TIME(*x*)

Displays *x* in a time format.

TIMESTAMP(*x*,*y*)

Creates a timestamp (date and time) from a date *x* and optionally a time *y*.

WEEK(*x*)

Determines the week of the year from date *x*.

YEAR(*x*)

Determines the year from date *x*.

Time and dates formats

Use any of the these formats when specifying times and dates.

Note: All date and time attributes in these rules are based in coordinated universal time (UTC).

Timestamp

Use one of the following formats to specify a timestamp:

- *date time*
- *date*

There must be exactly one space between the date and time.

You can mix formats for the date and time. For example, you can specify ISO format for the date and international format for the time.

Date Use one of these formats to specify a date:

European

DD.MM.YYYY

ISO *YYYY-MM-DD*

USA *MM/DD/YYYY*

You may leave off leading zeros from *MM* (month) and *DD* (day). You can use a two-digit year, in which case 1900 is added if the year is greater than 50 and 2000 is added if the year is 50 or less.

Note: The MONTHNAME() and DAYNAME() functions produce English names with no internationalization.

Time Use one of these formats to specify a time:

International

HH:MM[SS[.UUUUUU]]

USA *HH[:MM[:SS]] [A|P|AM|PM]*

You may leave off leading zeros from any field except subseconds. The international format uses a 24-hour clock. The USA format uses a 12-hour clock followed by A, P, AM, or PM.

You can substitute commas or periods for colon delimiters in the international format.

Examples

The following example shows a sample file.

```
VERSION 1

rule 'stgRule1' set stgpool 'pool1' for fileset ('cnt_A')
rule 'stgRule2' set stgpool 'pool2' where NAME like '%.doc'
rule 'stgRule3' set stgpool 'pool3' where DAYOFWEEK(CREATION_TIME) == 1
rule 'stgRule4' set stgpool 'pool4' where USER_ID <= 100
```

Creating a policy

This topic describes how to create a policy.

You must have Administrator privileges to perform this task.

SAN File System console provides a wizard to step you through the process of creating a policy.

Policy properties, including any associated rules, are stored in metadata. They are not stored in a file.

1. Start the Create-policy wizard by clicking **Manage Filing**→**Create a Policy** in the My Work frame.
2. Click **Next**.
3. In the Create a Policy panel under **High-Level settings**, fill in the Name and Description for the policy. Then click **Next**.
4. Under **Add Rules**, fill in the **Rules Description** field with a description of the rule.
5. Select a storage pool from the **Storage Pool Assignment** drop-down list.
6. Choose the rule specifics.
7. Continue creating all rules for this policy, and click **Next** when finished.
8. In the Edit Rules for Policy panel, verify the rules.
9. Click **Manage Filing**→**Policies**
10. Select the policy you just created, and click **Activate** from the drop-down list.
11. Click **Go** to activate the policy and verify the activation.

Sample policy sets

This topic provides sample policy sets.

Distribute files based on fileset

```
VERSION 1
RULE 'rule1' SET STGPOOL 'pool1' FOR FILESET('fileset1','fileset2')
RULE 'rule2' SET STGPOOL 'pool2' FOR FILESET('fileset3')
```

Distribute files based on file extension

```
VERSION 1
RULE 'documents' SET STGPOOL 'pool1' WHERE
    UCASE(NAME) LIKE '%.DOC' OR
    UCASE(NAME) LIKE '%.LWP' OR
    UCASE(NAME) LIKE '%.TXT'
RULE 'executables' SET STGPOOL 'pool2' WHERE
    UCASE(NAME) LIKE '%.EXE' OR
    UCASE(NAME) LIKE '%.COM' OR
    UCASE(NAME) LIKE '%.BAT' OR
    UCASE(NAME) LIKE '%.SH' OR
    UCASE(NAME) LIKE '%PL'
```

Distribute files based on the day of the week

Note:

1. The file placement resulting from this policy set cannot be restored from backups.
2. This policy set assumes placement based on coordinated universal time (UTC).

```
VERSION 1
RULE 'documents' SET STGPPOOL 'poo11' WHERE
  UCASE(NAME) LIKE '%.DOC' OR
  UCASE(NAME) LIKE '%.LWP' OR
  UCASE(NAME) LIKE '%.TXT'
RULE 'executables' SET STGPPOOL 'poo12' WHERE
  UCASE(NAME) LIKE '%.EXE' OR
  UCASE(NAME) LIKE '%.COM' OR
  UCASE(NAME) LIKE '%.BAT' OR
  UCASE(NAME) LIKE '%.SH' OR
  UCASE(NAME) LIKE '%PL'
```

Migrating data

This topic describes the procedures for migrating existing data to be managed by the SAN File System.

Data is migrated using the `migratedata` command from the client machine.

Attention: When you migrate journaled file system (JFS) files to SAN File System, you will lose access control lists (ACLs) from those files.

1. Estimate the time that it will take to migrate the data.
2. Import (or migrate) the data to the SAN File System.
3. Verify the integrity of the migrated data.

Estimating the time to migrate

This topic describes how to estimate the amount of time it takes to migrate data.

For large migrations, it is recommended that you estimate the amount of time that it will take to migrate the data set before you begin. The data-migration utility estimates this time based on several factors:

- Data-transfer rate over the storage area network (SAN)
- Amount of data being migrated
- Amount of available disk space on the target file system
- Amount of available memory
- Number of CPUs

To determine the data-transfer rate, the data-migration utility copies a set of the actual files from the source to the target file system.

Note: The estimation process can take a while if the data set is comprised of a large number of small files.

Review the data-migration prerequisites before you use the data-migration utility.

1. On the client machine, change to the directory where the `migratedata` command is located. For AIX, this is the `/usr/tank/migration/bin` directory. For Windows, this is the `c:\Program Files\IBM\Storage Tank\Migration` directory.
2. Invoke the `migratedata -phase plan` command.

Importing data into the SAN File System

This topic describes how to import data into the SAN File System.

Review the data-migration prerequisites before you begin migrating data.

Attention: When you migrate journaled file system (JFS) files to SAN File System, you will lose access control lists (ACLs) from those files.

You can migrate legacy data from your existing file system to the SAN File System using the data-migration utility on the client machine. This utility copies each file-system object from the source file system to the target SAN File System file system. The integrity of the migrated data and metadata (such as permissions and creation time) is checked automatically during the migration process.

The data-migration utility makes an entry in the log file before each file is migrated and marks that entry as "done" when the migration of that file is complete. When migrating large files, you can use the `-checkpoint` option to mark the entry in the log file after a specified number of blocks is migrated. The size of the block depends on the client platform.

Note: You can stop the data-migration process at any time and resume after the last completed file or block. The data-migration utility uses the log file to determine where the process was stopped; it knows where to resume the process.

1. On the client machine, change to the directory where the `migratedata` command is located. For AIX, this is the `/usr/tank/migration/bin` directory. For Windows, this is the `c:\Program Files\IBM\Storage Tank\Migration` directory.
2. Invoke the `migratedata -phase migrate` command.

Stopping a data migration

This topic describes how to stop a data migration currently in progress.

You can stop the data-migration process at any time and resume from the last completed file or block.

To stop the data-migration process, press `Ctrl+c`. You can then inspect the progress of the migration by viewing the log file.

Note: To clean up SAN File System after you stop the data migration, remove the migrated files using the standard operating system utilities, such as `rm` on UNIX or `del` on Windows, before you attempt to migrate data again.

Resuming a data migration

This topic describes how to resume a stopped data-migration process.

If you stop the data-migration process or if the process is terminated for some other reason, you can resume the migration after the last completed file or block without requiring a complete restart. If you migrated the same data set without specifying the `-restart` option again, the data-migration process starts from the beginning, and the data on the target file system will be overwritten.

To resume a data migration, invoke the `migratedata -phase migrate -resume` command on the SAN File System client machine.

Attention: You must specify the same log file as that used by the data-migration process being resumed. If you specify a different log file and do not specify the `-f` option, you will receive an error and the migration will stop. If you specify a different log file and specify the `-f` option, you will receive a warning and the data on the target file system will be overwritten.

Verifying the data integrity of migrated data

This topic describes how to verify migrated data.

The integrity of the migrated data and metadata (such as permissions and creation time) is checked automatically during the data-migration process.

You can also manually verify data integrity after the data migration is complete using either the data-migration utility or your own verification tools. The data-migration utility traverses both the source and target file systems and compares the metadata, file size, and checksum. Discrepancies in the attributes are reported and, if possible, repaired. Differences in file size or checksum are considered a failed migration. If a file appears in one file system but not in the other, the migration is also considered failed.

Review the data-migration prerequisites before you begin migrating data.

1. On the client machine, change to the directory where the `migratedata` command is located. For AIX, Linux, and Solaris, this is the `/usr/tank/migration/bin` directory. For Windows, this is the `c:\Program Files\IBM\Storage Tank\Migration` directory.
2. Invoke the `migratedata -phase verify` command.

Backing out migrated data

This topic describes how to back out migrated data.

You can back out a set of migrated data after the migration process is complete by pointing to the source file system rather than the SAN File System file system. The data-migration utility does not modify or delete data in the source file system.

Chapter 3. Uninstalling SAN File System

Uninstalling the package repository

This topic explains how to remove the SAN File System package repository.

Configuration settings and log files are not removed when you uninstall the package repository.

You must have root privileges to uninstall the package repository.

Use rpm to remove the package repository

```
rpm -e sfs-package
```

Uninstalling the metadata server

This topic explains how to remove the metadata server and the administrative server package.

Before uninstalling the metadata server, make sure that it is not running as part of the SAN File System cluster. In addition, you must have root privileges to uninstall the metadata server.

To view the packages that are installed, use the following command:

```
rpm -qa | grep sfs
```

1. Use rpm to remove the metadata server package

```
rpm -e sfs.server.linux
```

2. Use rpm to remove the administrative package

```
rpm -e sfs.admin.linux
```

Configuration settings and log files are not removed when you uninstall the metadata server.

Uninstalling the SAN File System software from a Windows client

This topic describes how to remove the SAN File System from a Windows client.

All applications that are currently running on the SAN File System client need to be stopped. For applications other than the SAN File System, refer to the documentation that comes with the application for information about stopping it.

1. From the Control Panel, double-click **Add/Remove Programs**.
2. Click **IBM SAN File System Client** in the Currently Installed Programs list.
3. Click **Change/Remove**.
4. Click **Yes**.
5. Click **Finish**.
6. Reboot the client.

Uninstalling the SAN File System software from an AIX client

This topic describes how to remove the SAN File System from an AIX client.

You must have root privileges to remove the SAN File System software from an AIX client.

1. Determine whether the client software is running.

```
mount
```

Determine if SAN File System is in the list of mounted file systems.

2. If the client is running, stop and unmount the SAN File System client for AIX.

```
/usr/tank/client/bin/rmstclient
```

This command unmounts the global namespace, stops the client, and unloads the kernel module. It uses the configuration information stored in `/usr/tank/client/config/stclient.conf`.

3. Use `installp` or `smit` to remove the software. For example:

```
installp -u . sfs.client.aix51
```

Uninstalling the SAN File System software from a Linux client

This topic provides describes how to remove the SAN File System from a Linux client.

You must have root privileges to remove the SAN File System software from a Linux client.

1. Determine whether the client software is running.

```
mount
```

Determine if SAN File System is in the list of mounted file systems.

2. If the client is running, stop and unmount the SAN File System client for Linux.

```
/usr/tank/client/bin/rmstclient
```

This command unmounts the global namespace, stops the client, and unloads the kernel module. It uses the configuration information stored in `/usr/tank/client/config/stclient.conf`.

3. Use **rpm** to remove the software. For example:

```
rpm -e sfs.client.linux_RHEL3_9
```

Uninstalling the SAN File System software from a Solaris client

This topic provides describes how to remove the SAN File System from a Solaris client.

You must have root privileges to remove the SAN File System software from a Solaris client.

1. Determine whether the client software is running.

```
mount
```

Determine if SAN File System is in the list of mounted file systems.

2. If the client is running, stop and unmount the SAN File System client for Solaris.

```
/usr/tank/client/bin/rmstclient
```

This command unmounts the global namespace, stops the client, and unloads the kernel module. It uses the configuration information stored in `/usr/tank/client/config/stclient.conf`.

3. Remove the SAN File System software from the client using the Solaris package program.

```
pkgrm IBMSFS
```

Chapter 4. Upgrading the SAN File System from Version 1.1

This topic provides an overview of the procedure to upgrade the SAN File System from Version 1.1 to Version 2.1.

Before you begin upgrading the SAN File System, make sure that you have access to all of the required software. Most of the prerequisite software that you will need is available on the SAN File System CD-ROM.

You will need to obtain the following software:

- SUSE LINUX Enterprise Server 8.0. You will need a licensed copy of LINUX Enterprise Server 8.0 for each of the metadata server engines in the cluster. For more information about obtaining Enterprise Server, visit www.suse.com.
- Qlogic driver. For information about obtaining the QLogic driver, visit http://www.qlogic.com/support/oem_detail_all.asp?oemid=22.
- United Linux Service Pack 3. For more information about obtaining the United Linux Service Pack 3, visit www.suse.com.

Use the following checklist to upgrade the SAN File System from version 1.1:

Steps		For more information...
1	Verify the current state of your SAN File System.	"Verifying the current state of the SAN File System" on page 77
2	Using the SAN File System version 1 documentation, create a complete backup of your existing system.	
3	Upgrade the master console.	<i>Master Console User's Guide</i>
4	Upgrade the first metadata server engine.	"Upgrading the first metadata server engine" on page 80
a	Prepare the engine for upgrade.	"Preparing the first engine for upgrade" on page 80
1	Obtain software prerequisites.	"Obtain prerequisite software" on page 30
2	Disable the high-availability script.	"Disable the high-availability script" on page 80
3	Disable RSA II timer watchdogs.	"Disabling the RSA II watchdogs" on page 81
4	Optionally, reassign filesets to other engines.	"Reassign filesets to another metadata server in the cluster" on page 81
5	Copy configuration files to another location.	"Copy configuration files to a backup location" on page 82
6	Upgrade system BIOS.	"Upgrading system BIOS" on page 30
7	Upgrade RSA II firmware.	"Upgrading RSA II firmware" on page 42
8	Optionally, split the boot drive mirror.	"Splitting the boot drive mirror" on page 83

Steps		For more information...
b	Install all software.	"Installing software on a metadata server" on page 84
1	Install SUSE LINUX Enterprise Server 8.0.	"Installing the operating system" on page 31
2	Disable the X Window System.	"Disabling the automatic starting of the X Window System" on page 34
3	Set the date and time on the engine.	"Setting the time and date on the Metadata servers" on page 34
4	Apply United Linux Service Pack 3 (SP3) updates.	"Apply United Linux Service Pack 3 (SP3) updates" on page 34
5	Install prerequisite software.	"Installing prerequisite software" on page 35
a	QLogic driver.	"Install QLogic driver" on page 35
b	MPCLI.	"Install MPCLI" on page 38
c	Java runtime environment.	"Install the Java Runtime Environment" on page 38
d	IBM Director.	"Install the IBM Director agent" on page 38
e	Eclipse.	"Install Eclipse" on page 38
f	Ibmusbasm.	"Install ibmusbasm" on page 38
g	Openslp.	"Install OpenSLP" on page 39
h	IBM SDD.	"Install the IBM Subsystem Device Driver (SDD)" on page 39
i	WebSphere 5.0 Express.	"Install IBM WebSphere 5.0 Express" on page 39
6	Install the SAN File System software.	"Installing SAN File System software" on page 40
7	Restore configuration files.	"Restore configuration settings" on page 84
c	Set up the first metadata server engine.	"Setting up the metadata server" on page 85
d	Re-enable the boot drive mirror.	"Re-enabling the boot drive mirror" on page 85
e	Set the new master metadata server.	"Setting the upgraded server as the new master" on page 87
1	Transfer the master role to the upgraded engine.	"Transferring the master role to the upgraded metadata server" on page 87
2	Optionally, reassign filesets to new master.	"Reassign filesets to another metadata server in the cluster" on page 88
3	Configure SNMP traps on the new master.	"Configuring metadata servers for SNMP traps" on page 54

Steps		For more information...
	4 Perform a backup of the new master. Important: After upgrading the first metadata server and making it the new master, you will need to use version 2.1 command syntax when you issue SAN File System administrative commands from the master metadata server.	"Performing a backup of the new master" on page 90
5	Upgrade each remaining subordinate in the cluster.	"Upgrading additional metadata servers" on page 91
	a Prepare the engine for upgrade.	"Preparing an engine for upgrade" on page 91
	1 Obtain software prerequisites.	"Obtain prerequisite software" on page 30
	2 Stop the metadata server.	"Stopping a metadata server" on page 81
	3 Disable RSA II timer watchdogs.	"Disabling the RSA II watchdogs" on page 81
	4 Optionally, reassign filesets to other engines.	"Reassign filesets to another metadata server in the cluster" on page 81
	5 Copy configuration files to another location.	"Copy configuration files to a backup location" on page 82
	6 Upgrade system BIOS.	"Upgrading system BIOS" on page 30
	7 Upgrade RSA II firmware.	"Upgrading RSA II firmware" on page 42
	8 Optionally, split the boot drive mirror.	"Splitting the boot drive mirror" on page 83
	b Install all software.	"Installing software on a metadata server" on page 92
	1 Install SUSE LINUX Enterprise Server 8.0.	"Installing the operating system" on page 31
	2 Disable the X Window System.	"Disabling the automatic starting of the X Window System" on page 34
	3 Set the date and time on the engine.	"Setting the time and date on the Metadata servers" on page 34
	4 Apply United Linux Service Pack 3 (SP3) updates.	"Apply United Linux Service Pack 3 (SP3) updates" on page 34
	5 Install prerequisite software.	"Installing prerequisite software" on page 35
	a QLogic driver.	"Install QLogic driver" on page 35
	b MPCLI.	"Install MPCLI" on page 38
	c Java runtime environment.	"Install the Java Runtime Environment" on page 38
	d IBM Director.	"Install the IBM Director agent" on page 38
	e Eclipse.	"Install Eclipse" on page 38
	f Ibmusbasm.	"Install ibmusbasm" on page 38

Steps		For more information...
	g Openslp.	"Install OpenSLP" on page 39
	h IBM SDD.	"Install the IBM Subsystem Device Driver (SDD)" on page 39
	i WebSphere 5.0 Express.	"Install IBM WebSphere 5.0 Express" on page 39
6	Install the SAN File System software.	"Installing software on a metadata server" on page 92
7	Restore configuration files.	"Restore configuration settings" on page 84
c	Set up the metadata server engine.	"Setting up a metadata server engine" on page 92
d	Verifying the installation.	"Verifying the installation" on page 89
1	Optionally, reassign filesets to the upgraded metadata server.	"Reassign filesets to another metadata server in the cluster" on page 88
2	Configure SNMP traps on the upgraded metadata server.	"Configuring metadata servers for SNMP traps" on page 54
6	Upgrade clients.	
a	Upgrade Windows clients.	"Upgrading SAN File System on a Windows client" on page 93
1	Optionally, install SDD, RDAC, or other multipathing software.	For information about installing SDD, see "Installing SDD on clients" on page 47
2	Obtain Windows client software.	"Obtain version 2.1 software for a Windows client" on page 47
3	Prepare the client for upgrade.	"Preparing a Windows client for upgrading" on page 94
a	Stop all client applications.	"Stop all applications on the client" on page 94
b	Uninstall the Windows client.	"Uninstalling the SAN File System software from a Windows client" on page 69
4	Install the software.	"Installing the SAN File System software" on page 48
5	Validate client installation.	"Validating the installation of SAN File System on a Windows client" on page 48
6	Set up the client to automatically start when the client is rebooted.	"Automate client restart on reboot" on page 94
b	Upgrade AIX clients.	"Upgrading SAN File System on an AIX client" on page 95
1	Optionally, install SDD, RDAC, or other multipathing software.	For information about installing SDD, see "Installing SDD on clients" on page 47
2	Obtain AIX client software.	"Obtain version 2.1 software for an AIX client" on page 49
3	Prepare the client for upgrade.	"Preparing an AIX client for upgrading" on page 95

Steps		For more information...
	a Stop all client applications.	"Stop all applications on the client" on page 96
4	Install the software.	"Installing the SAN File System software" on page 50
5	Validate client installation.	"Validating the installation of SAN File System on an AIX client" on page 51
7	Complete the upgrade process.	
	a Verify new installation.	"Verifying the installation" on page 97
	b Upgrade cluster.	"Upgrading the cluster" on page 97
	c Revise scripts and policies.	"Revising scripts and policies" on page 98
	d Modify zoning.	"Modifying metadata server zoning and LUN masking" on page 99
8	Back up the complete system.	Chapter 6, "Backing up the SAN File System," on page 111

Verifying the current state of the SAN File System

This topic provides the procedure for verifying that the engine is installed.

Before you begin the upgrade process, you can run several administrative command-line interface commands and save the results to a file. You can then compare the output of these commands before and after the upgrade to ensure the SAN File System is functioning as you expect.

When you save the configuration files (see "Copy configuration files to a backup location" on page 82), make sure that you save these files as well.

Run these commands from the master metadata server engine.

1. If SDD is installed:

a. Run the **lsvpcfg** command.

```
lsvpcfg > lsvpcfg_v1
```

The file should contain results similar to the following:

```
000 vpathai ( 254, 128) 20617844 = /dev/sdb /dev/sdg
001 vpatham ( 254, 192) 30017844 = /dev/sdc /dev/sdh
002 vpathan ( 254, 208) 30117844 = /dev/sdd /dev/sdi
003 vpathao ( 254, 224) 30217844 = /dev/sde /dev/sdj
004 vpathap ( 254, 240) 30317844 = /dev/sdf /dev/sdk.
```

b. Run the **datapath query device** command.

```
datapath query device > datapath_v1
```

The file should contain results similar to the following:

```
Total Devices : 5
```

```
DEV#: 0 DEVICE NAME: vpathai TYPE: 2105F20 POLICY: Optimized
SERIAL: 20617844
```

```
=====
Path#      Adapter/Hard Disk      State      Mode      Select      Errors
```

```

0      Host2Channel0/sdb      OPEN  NORMAL      488848      0
1      Host3Channel0/sdg      OPEN  NORMAL      490227      0

```

```

DEV#: 1  DEVICE NAME: vpatham  TYPE: 2105F20  POLICY: Optimized
SERIAL: 30017844

```

```

=====
Path#   Adapter/Hard Disk   State   Mode       Select   Errors
0       Host2Channel0/sdc   OPEN    NORMAL     251670   0
1       Host3Channel0/sdh   OPEN    NORMAL     251434   0

```

```

DEV#: 2  DEVICE NAME: vpathan  TYPE: 2105F20  POLICY: Optimized
SERIAL: 30117844

```

```

=====
Path#   Adapter/Hard Disk   State   Mode       Select   Errors
0       Host2Channel0/sdd   CLOSE   NORMAL     1480     0
1       Host3Channel0/sdi   CLOSE   NORMAL     1542     0

```

```

DEV#: 3  DEVICE NAME: vpathao  TYPE: 2105F20  POLICY: Optimized
SERIAL: 30217844

```

```

=====
Path#   Adapter/Hard Disk   State   Mode       Select   Errors
0       Host2Channel0/sde   CLOSE   NORMAL     1492     0
1       Host3Channel0/sdj   CLOSE   NORMAL     1530     0

```

```

DEV#: 4  DEVICE NAME: vpathap  TYPE: 2105F20  POLICY: Optimized
SERIAL: 30317844

```

```

=====
Path#   Adapter/Hard Disk   State   Mode       Select   Errors
0       Host2Channel0/sdf   CLOSE   NORMAL     1550     0
1       Host3Channel0/sdk   CLOSE   NORMAL     1472     0

```

2. Run the **lsserver** command.

```
tanktool lsserver > lsserver_v1
```

The file should contain results similar to the following:

```

Name      State  Server Role Containers Last Boot
=====
ivt1-mds1 Online Master                3 Mar 09, 2004 2:49:09 PM
ivt1-mds2 Online Subordinate    2 Mar 09, 2004 2:40:57 PM

```

3. Run the **lsadmuser** command.

```
tanktool lsadmuser > lsadmuser_v1
```

The file should contain results similar to the following:

```

Name      User Role Authorization
=====
monitor   Monitor Not Current
admin     Admin  Not Current
MDMSUID   Admin  Not Current
operator  Operator Not Current
backup    Backup Not Current
itpc      Admin  Not Current
superuser Admin  Not Current
test      Admin  Current

```

4. Run the **lsautorestart** command.

```
tanktool lsautorestart > lsautorestart_v1
```

The file should contain results similar to the following:

```

Name      Service State Probe State Last Probe      Probes Highest Retries
=====
ivt1-mds2 Running      Live Server Apr 06, 2004 2:49:34 PM  7477      0
ivt1-mds1 Running      Live Server Apr 06, 2004 2:57:13 PM  6658      0

```

5. Run the **lsclient** command.

```
tanktool lsclient > lsclient_v1
```

The file should contain results similar to the following:

Client	Session ID	State	Server	Renewals	Privilege
ivt1-c2-win	8	Current	ivt1-mds1	257075	Root
ivt1-c3-aix	3	Current	ivt1-mds1	331199	Root
ivt1-c1-win	1	Current	ivt1-mds1	331079	Root
ivt1-c2-win	9	Current	ivt1-mds2	257055	Root
ivt1-c3-aix	4	Current	ivt1-mds2	331200	Root
ivt1-c1-win	2	Current	ivt1-mds2	331096	Root

6. Run the **lscontainer** command.

```
tanktool lscontainer > lscontainer_v1
```

The file should contain results similar to the following:

Name	Container	State	Quota Type	Quota (MB)	Used (MB)	Used (%)	Threshold (%)	Most...
ROOT	Attached	Soft		0	16	0	0	-
CONTAINER1	Attached	Soft		0	3984	0	80	-
CONTAINER2	Attached	Soft		0	3952	0	80	-
CONTAINER3	Attached	Soft		0	3920	0	80	-
CONTAINER4	Attached	Soft		0	3904	0	80	-

7. Run the **lslun** command.

```
tanktool lslun > lslun_v1
```

The file should contain results similar to the following:

OS	Device	Path	Lun ID	Storage Device	WWNN	Vendor	Product	Size (MB)	Volume	State
/dev/rvpatha	20027074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_00	Assigned			
/dev/rvpathb	20127074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_01	Assigned			
/dev/rvpathc	20227074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_02	Assigned			
/dev/rvpathd	20327074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_03	Assigned			
/dev/rvpathe	20427074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_04	Assigned			
/dev/rvpathf	20527074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_05	Assigned			
/dev/rvpathg	20627074	50:05:07:63:00:c4:a7:5a	IBM	2105800	9536	vol_06	Assigned			

8. Run the **lspool** command.

```
tanktool lspool > lspool_v1
```

The file should contain results similar to the following:

Name	Type	Size (MB)	Used (MB)	Used (%)	Threshold (%)	Volumes
SYSTEM	System	71520	1616	2	80	5
DEFAULT_POOL	User Default	9536	16	0	80	1
TEST_POOL1	User	143040	7936	5	90	12
TEST_POOL2	User	157344	7824	4	90	12

9. Run the **lsvol** command.

```
tanktool lsvol > lsvol_v1
```

The file should contain results similar to the following:

Name	State	Pool	Size (MB)	Used (MB)	Used (%)
MASTER	Activated	SYSTEM	14304	336	2
SYSTEM1	Activated	SYSTEM	14304	336	2
SYSTEM2	Activated	SYSTEM	14304	320	2
SYSTEM3	Activated	SYSTEM	14304	320	2
SYSTEM4	Activated	SYSTEM	14304	304	2
vol_00	Activated	DEFAULT_POOL	9536	16	0
vol_01	Activated	TEST_POOL1	9536	656	6
vol_02	Activated	TEST_POOL1	9536	640	6
vol_03	Activated	TEST_POOL1	9536	640	6
vol_04	Activated	TEST_POOL1	9536	656	6

Upgrading the first metadata server engine

This topic provides the general steps for upgrading the first metadata server engine.

Choose an existing subordinate metadata server to be the first metadata server to be upgraded. After you have completed the upgrade process and the metadata server has rejoined the cluster, you will transfer the master role to this metadata server.

1. Prepare the metadata server engine for upgrade so that the SAN File System cluster continues to function while you are upgrading this metadata server engine. See “Preparing the first engine for upgrade.”
2. Install all software on the metadata server engine. See “Installing software on a metadata server” on page 84.
3. Restore the configuration settings. See “Restore configuration settings” on page 84.
4. Set up the metadata server engine and validate that it is installed correctly by having it rejoin the cluster and ensuring that this metadata server engine has access to LUNs and volumes. See “Setting up the metadata server” on page 85.
5. If you disabled the boot drive mirror, you can now re-enable the mirror. See “Re-enabling the boot drive mirror” on page 85.
6. Set this metadata server engine as the new master, which includes performing a backup of the new configuration. See “Setting the upgraded server as the new master” on page 87.

Preparing the first engine for upgrade

This topic describes how to prepare the first engine to be upgraded.

1. Obtain software to be used on the upgrade process. See “Obtain prerequisite software” on page 30.
2. Disable the high-availability script. See “Disable the high-availability script.”
3. Stop the engine. See “Stopping a metadata server” on page 81.
4. Disable RSA II timer watchdogs. See “Disabling the RSA II watchdogs” on page 81
5. Optionally, reassign filesets to other engines. See “Reassign filesets to another metadata server in the cluster” on page 81
6. Copy configuration files to another location. See “Copy configuration files to a backup location” on page 82.
7. Upgrade the system BIOS. See “Upgrading system BIOS” on page 30.
8. Upgrade the RSA II firmware. See “Upgrading RSA II firmware” on page 42.
9. Optionally, split the boot drive mirror. See “Splitting the boot drive mirror” on page 83.

Disable the high-availability script

This topic describes how to disable the high-availability script.

The high-availability script is used to automatically transfer workload from a failing metadata server engine to another engine in the cluster. Therefore, you will need to disable the high-availability script before you begin upgrading to version 2.1.

Note: You only need to disable the high-availability script one time, from the master metadata server.

Disable the high-availability script if it is running.

```
/usr/tank/admin/bin/tanktool legacy "setautofailover disable"
```

Stopping a metadata server

This topic describes how to stop a metadata server engine.

From the master metadata server, stop the metadata server to be upgraded.

```
/usr/tank/admin/bin # tanktool stopserver metadata_server_name
```

Disabling the RSA II watchdogs

This topic describes how to disable the RSA II watchdogs.

You need to disable the RSA II adapter watchdogs before you begin upgrading the engine. Otherwise, the RSA II adapter may attempt to automatically restart the engine when you power it off.

1. Make sure the engine is powered on.
2. From the master console (or any Windows client with network access to the RSA II adapter), open a Web browser and point it to the IP address for the RSA II adapter.
3. Log on to the RSA II adapter.
4. In the left frame, click **Server --> ASM Control --> System Settings**.
5. Under Server Timeouts, set these watchdogs to **Disabled**.
 - POST watchdog
 - OS watchdog
 - Loader watchdog
6. Go to the bottom of the page and click **Save** to save your settings.
7. Under ASM control, click **Restart ASM** to enable your changes.
8. From the Restart ASM panel, click **Restart**.
9. When prompted to confirm that you want to restart the adapter, click **OK**.
10. When prompted to close the window, click **Yes**.

Reassign filesets to another metadata server in the cluster

This topic describes how to reassign filesets to another metadata server in the SAN File System cluster.

Any filesets managed by this metadata server engine will be unavailable to clients while the engine is being upgraded. If it is important to maintain the availability of the filesets during the upgrade process, you will need to reassign the filesets to another metadata server in the cluster.

Note: Before reassigning filesets, you should close down any client applications that use those filesets until after the reassignment is complete.

1. From the master metadata server, list all of the filesets assigned to the metadata server to be upgraded.

```
/usr/tank/admin/bin/tanktool lscontainer -server metadata_server_name
```

2. Assign the filesets to different metadata server.

```
/usr/tank/admin/bin/tanktool setcontainerserver -server new_metadata_server_name  
fileset1 fileset2 fileset3
```

There are two phases of the cluster transition when moving a fileset. The end of the first phase shows the server online, but some administrative commands will not function until the second phase completes. The second phase finishes when the new filesets are assigned and running. This normally takes less than one minute.

Copy configuration files to a backup location

This topic describes how to copy the networking and SAN File System configuration files before upgrading the SAN File System.

The network and SAN File System configuration files that are currently on the metadata server engine are needed to restart the metadata server engine and add it back to the cluster after upgrading the engine. However, upgrading the metadata server engine to version 2.1 requires that you install a new operating system. Therefore, you will need to tar and copy the configuration files to another location before you proceed.

You can copy the files to any location to which you will have access after upgrading the metadata server. For example, you can copy the configuration files to a diskette.

The following network configuration files need to be included:

- /etc/HOSTNAME
- /etc/sysconfig/network/ifcfg-eth*
- /etc/sysconfig/network/routes
- /etc/hosts
- /etc/resolv.conf

The following SAN File System configuration files need to be included:

- /root/.tank.passwd
 - /etc/password (if it exists)
 - /etc/shadow (if it exists)
 - /etc/tank/server/Tank.Bootstrap
 - /etc/tank/server/Tank.Config
 - /etc/tank/admin/config/tank.properties
 - /etc/tank/admin/config/cimom.properties
 - /usr/tank/admin/truststore
 - /var/tank/server/DR/
 - All of the version 1 administrative command-line interface files you created when you verified the current state of the SAN File System. See “Verifying the current state of the SAN File System” on page 77 for more information.
1. From the metadata server to be upgraded, tar the network configuration files into a single file. For example, to create a file called `network_config_mds2`:

```
tar -cPvf network_config_mds2 /etc/HOSTNAME /etc/sysconfig/network/ifcfg-eth* /etc/sysconfig/network/routes /etc/hosts /etc/resolv.conf
```
 2. Create a tar file containing the SAN File System configuration files into a single file. For example, to create a file called `sfs_config_mds2`:

Remember: Do not forget to include the version 1 files you created when you verified the current state of the SAN File System.

```
tar -cPvf network_config_mds2 /root/.tank.passwd /etc/tank/server/Tank.Bootstrap /etc/tank/server/Tank.Config /etc/tank/admin/config/tank.properties ...remaining files ...
```

3. Copy the configuration tar file to a diskette.

Tip: Rather than copy the file to diskette, you can copy the file to another location in the network, such as the master console. To copy files to the

master console, use the **pscsp** and **pftp** commands. For information about using these commands, see the PuTTY documentation that is available on the master console.

If you copy the files to another location in the network, you will need to ensure that you have network connectivity to restore them when needed.

- a. Mount the diskette drive.

```
mount /dev/fd0 /media/floppy
```

- b. Copy the network tar file to a diskette.

```
cp network_config_mds2 /media/floppy
```

- c. Copy the SAN File system configuration tar file to a diskette.

```
cp sfs_config_mds2 /media/floppy
```

- d. Verify that the tar files are on the diskette.

```
ls /media/floppy
```

- e. Unmount the diskette drive and remove the diskette.

```
umount /media/floppy
```

Splitting the boot drive mirror

This topic describes the steps for splitting the mirrored boot drives into two separate drives.

You should consider temporarily disabling the mirroring capabilities of the boot drive before you begin upgrading to version 2.1. This allows you to keep version 1.1 of the SAN File System on one drive while you upgrade to version 2.1 using the other drive. If you need to restore your version 1.1 boot drive configuration, you can reconfigure your boot drive mirror with the secondary drive as the “primary mirror device,” and your recently installed primary drive as the “secondary mirror device.” Otherwise, after you are satisfied with your upgrade to version 2.1, you can recreate the mirror with the newly installed primary drive as the “primary mirror device” and the secondary drive as the “secondary mirror device.”

Note: Keep in mind that while the mirrored drives are split, the high availability protection provided by the RAID mirroring is disabled (each drive becomes a single point of failure during this installation window).

For information about splitting the mirrored boot drive, see the *Device Management Users Guide*, which is available at this website:

www.lsillogic.com

1. Reboot the engine.
2. Watch the boot sequence carefully and press Ctrl-c when the LSI message appears.
3. Select the first device and press Enter.
4. Select **Mirroring Properties**.
5. Change the properties from Primary and Secondary to **No**.
6. Save the configuration and exit.
7. After the engine finishing booting, perform a test mount to verify that you have access to the boot drive.

```
mount /dev/sdb1 /mnt
umount /mnt
```

Installing software on a metadata server

This topic provides an overview of the steps required to install the software on a metadata server to be upgraded.

The steps for installing the software are the same as the steps for installing software on the first metadata server engine.

1. Install SUSE LINUX Enterprise Server 8.0. See “Installing the operating system” on page 31
2. Disable the automatic starting of X-Windows (if you choose to have graphical mode as the default desktop setting during the installation of the operating system). See “Disabling the automatic starting of the X Window System” on page 34.
3. Set the time and date on the metadata server engine. See “Setting the time and date on the Metadata servers” on page 34.
4. Apply the United Linux Service Pack 3. See “Apply United Linux Service Pack 3 (SP3) updates” on page 34
5. Restore configuration settings from your backup location. See “Restore configuration settings”
6. Install all prerequisite software, such as the QLogic driver and the IBM Director agent. See “Installing prerequisite software” on page 35
7. Install the SAN File System software. See “Installing SAN File System software” on page 40

Restore configuration settings

You need to restore the configuration settings for your network.

You need to restore the following networking configuration files from your backup location:

- /etc/HOSTNAME
- /etc/sysconfig/network/ifcfg-eth*
- /etc/sysconfig/network/routes
- /etc/hosts
- /etc/resolv.conf

In addition, you need to restore the following SAN File System configuration files from your backup location:

- /root/.tank.passwd
- /etc/password (if it exists)
- /etc/shadow (if it exists)
- /etc/tank/server/Tank.Bootstrap
- /etc/tank/server/Tank.Config
- /etc/tank/admin/config/tank.properties
- /etc/tank/admin/config/cimom.properties
- /usr/tank/admin/truststore
- /var/tank/server/DR/
- All of the version 1 administrative command-line interface files you created when you verified the current state of the SAN File System. See “Verifying the current state of the SAN File System” on page 77 for more information.

If you followed the steps for “Copy configuration files to a backup location” on page 82, you can use the following procedure to restore those files to the metadata server engine.

1. Insert the diskette into the diskette drive and mount the diskette drive.

Tip: If you copied the tar files to the master console instead of putting them on a diskette, use the **pscsp** and **pftp** commands to restore them. For information about using these commands, see the PuTTY documentation that is available on the master console.

```
mount /dev/fd0 /media/floppy
```

2. Change directories to the diskette drive.

```
cd /media/floppy
```

3. Untar the network configuration files

```
tar -xPvf network_config_mds2.tar
```

4. Untar the SAN File System configuration files

```
tar -xPvf sfs_config_mds2.tar
```

5. Change directories back to the root directory.

```
cd /root
```

6. Unmount the diskette drive and remove the diskette.

```
umount /media/floppy
```

7. Reboot the server to apply the changed configuration files.

Setting up the metadata server

This topic describes how to set up the first engine after you have installed all software.

1. Ensure that the engine is powered ON and that you are logged in as root.
2. Start the metadata server.

```
/usr/tank/admin/bin/setupfs -noprompt
```

3. From the master metadata server, list all servers in the SAN File system cluster to verify that the metadata server has rejoined the cluster.

```
/usr/tank/admin/bin/tanktool lserver
```

Note: At this point, the current master metadata server is still running in version 1 mode. Therefore, you invoke the ACLI using `tanktool`.

However, after you have transferred the master role to the upgraded metadata server, you will invoke the ACLI using `sfscli`.

You should see the upgraded server in the list of metadata servers:

```
Name      State  Server Role Containers Last Boot
=====
mstr_mds  Online Master                3 Feb 13, 2004 2:46:28 PM
sub_mds1  Online Subordinate           1 Feb 16, 2004 4:26:08 AM
```

If you did not move the workload from the metadata server before upgrading, the filesets managed by the upgraded metadata server should again be accessible to SAN File System client.

Re-enabling the boot drive mirror

This topic describes the steps for to re-enable mirroring on the boot drives.

If you split the boot drive mirror before you began the upgrade, you can now re-enable the boot drive mirror.

For information about splitting the mirrored boot drive, see the *Device Management Users Guide*, which is available at this website:

www.lsilogic.com

1. From the master metadata server, stop the metadata server that was just updated.

```
/usr/tank/admin/bin/tanktool stopserver metadata_server_name
```

2. Exclude the secondary drive from the list of drives that are available.

- a. Edit `/etc/fstab`.

```
/dev/sda3      /          reiserfs  defaults          1 1
/dev/sda1      /data1     auto      noauto,user       0 0
/dev/sdb1      /data2     auto      noauto,user       0 0
/dev/sdb3      /data3     auto      noauto,user       0 0
/dev/sda2      swap       swap      pri=42            0 0
/dev/sdb2      swap       swap      pri=42            0 0
devpts        /dev/pts   devpts    mode=0620,gid=5   0 0
proc          /proc      proc      defaults          0 0
usbdevfs      /proc/bus/usb usbdevfs  noauto            0 0
/dev/cdrom     /media/cdrom auto      ro,noauto,user,exec 0 0
/dev/fd0       /media/floppy auto      noauto,user,sync  0 0
```

- b. Delete all `/dev/sdb` partitions.

Important: You must actually delete the `/dev/sdb` partitions from this file rather than commenting them out.

```
/dev/sda3      /          reiserfs  defaults          1 1
/dev/sda1      /data1     auto      noauto,user       0 0
/dev/sda2      swap       swap      pri=42            0 0
devpts        /dev/pts   devpts    mode=0620,gid=5   0 0
proc          /proc      proc      defaults          0 0
usbdevfs      /proc/bus/usb usbdevfs  noauto            0 0
/dev/cdrom     /media/cdrom auto      ro,noauto,user,exec 0 0
/dev/fd0       /media/floppy auto      noauto,user,sync  0 0
```

- c. Save the file.

3. Reboot the engine to access the LSI configuration utility.
4. Watch the boot sequence carefully, and press Ctrl-c when the LSI message appears.
5. Select the first device and press Enter.
6. Select **Mirroring Properties**.
7. Change the properties from No to **Primary and Secondary**.
8. Save the configuration and exit.
9. After the engine finishing booting, perform a test mount to verify that you have access to the boot drive.

```
mount /dev/sdb1 /mnt
```

10. Unmount the drive.

11. Start the metadata server.

```
/usr/tank/admin/bin/tanktool startserver metadata_server_name
```

12. From the master metadata server, list all servers in the SAN File system cluster to verify that the metadata server has rejoined the cluster.

```
/usr/tank/admin/bin/tanktool lserver
```

You should see the upgraded server in the list of metadata servers:

Name	State	Server Role	Containers	Last Boot
mstr_mds	Online	Master		3 Feb 13, 2004 2:46:28 PM
sub_mds1	Online	Subordinate		1 Feb 16, 2004 4:26:08 AM

Setting the upgraded server as the new master

This topic provides an overview of the steps required to set the metadata server that you have just upgraded as the new master metadata server for the cluster.

1. Transfer the master role to the upgraded metadata server. See “Transferring the master role to the upgraded metadata server.”
2. Optionally, reassign filesets to the new master. See “Reassign filesets to another metadata server in the cluster” on page 88.
3. Configure SNMP traps on the new master. See “Configuring metadata servers for SNMP traps” on page 54.
4. Back up the new master configuration. See “Performing a backup of the new master” on page 90.

Transferring the master role to the upgraded metadata server

This topic describes how to transfer the master role to the upgraded metadata server engine.

If this metadata server is the first metadata server to be upgraded, you will need to make it the new master metadata server.

Note: Transferring the master role to the upgraded metadata server requires that you shut down the current master and then make this metadata server the new master. During this procedure, the SAN File System will be unavailable to all clients.

1. Optionally, move the current master’s workload to one of the subordinate nodes. By moving the workload, the filesets being managed by the current master will be unavailable only while you are transferring the master role. Otherwise, the filesets will be unavailable until you have upgraded the current master.
2. From the master metadata server, take the current master offline.

```
/usr/tank/admin/bin/tanktool stopserver master_mds_name
```
3. Verify that the current master is offline

```
/usr/tank/admin/bin/tanktool lsserver
```

The status of the current master should be listed as Not Running.

```
/usr/tank/admin/bin/tanktool lsserver
Name      State      Server Role Filesets Last Boot
=====
evt5-mds1 Not Running Master      -      -
evt5-mds2 Joining   Subordinate      0 Feb 27, 2004 5:18:27 AM
```

4. From the upgraded metadata server, verify the cluster status.

```
/usr/tank/admin/bin/sfsccli lsserver
```

Because you are running the lsserver command from a subordinate, you will only see the status of the upgraded metadata server.

```
evt5-mds2:/usr/tank/admin/bin/sfsccli lsserver
Name      State      Server Role Filesets Last Boot
=====
evt5-mds2 Joining   Subordinate      0 Feb 27, 2004 5:18:27 AM
```

5. Make the upgraded metadata server the new master

```
/usr/tank/admin/bin/sfsccli setmaster upgraded_mds_name
```

6. From the new master metadata server, verify that the cluster has been reformed properly.

```
/usr/tank/admin/bin/sfsccli lsserver
```

The upgraded metadata server should now be listed as the master. All other metadata servers should be listed as subordinates with a status of active.

Note: The former master metadata server will still show a status of Not Running.

```
Name      State      Server Role Filesets Last Boot
=====
evt5-mds2 Online      Master      2 Feb 16, 2004 6:15:39 AM
evt5-mds1 Not Running Subordinate - -
```

7. Verify the software version and committed software version of the cluster:

```
/usr/tank/admin/bin/sfsccli statcluster
```

The software version should be at version 2.1 (such as 2.1.0.65), but the committed software version should still be at version 1.1 (such as 1.1.2.8).

```
Name                sanfs
ID                  21443
State               Online
Target State       Online
Last State Change   Feb 16, 2004 6:15:47 AM
Last Target State Change -
Servers             2
Active Servers      1
Software Version    2.1.0.65
Committed Software Version 1.1.2.8
Last Software Commit Feb 13, 2004 2:46:27 PM
Software Commit Status Not In Progress
Installation Date    Feb 13, 2004 2:46:27 PM
```

The SAN File System is now back online and fully operational (with the exception of any filesets that are still being managed by the former master metadata server). At this point, the SAN File System cluster is running in version 1-compatibility mode. This means that SAN File System clients and metadata servers that have not been upgraded will still continue to operate. However, they will not take advantage of any version 2.1 functionality.

Note: On metadata servers that have been upgraded to version 2.1, you will need to use the version 2.1 syntax for all administrative commands. On metadata servers that have not been upgraded, you will need to use version 1 syntax for all administrative commands.

Reassign filesets to another metadata server in the cluster

This topic describes how to reassign filesets to another metadata server in the SAN File System cluster from Version 2.1.

Any filesets managed by this metadata server engine will be unavailable to clients while the engine is being upgraded. If it is important to maintain the availability of the filesets during the upgrade process, you will need to reassign the filesets to another metadata server in the cluster.

The master metadata server has already been upgraded. Therefore, you will need to run the Version 2.1 commands to reassign the filesets.

1. List all of the filesets assigned to the metadata server to be upgraded.

```
/usr/tank/admin/bin/sfsccli lsfileset -server metadata_server_name
```

2. Assign the filesets to different metadata server.

```
/usr/tank/admin/bin/sfsccli setfilesetserver -server new_metadata_server_name  
fileset1 fileset2 fileset3
```

Verifying the installation

This topic provides the procedure for verifying that the engine is installed.

You need to run the administrative command-line interface commands and save them to a file. You can then compare them with the version 1 commands you ran to ensure that the system is functioning normally.

Note: Keep in mind that the format of the commands may change slightly in version 2.1 of the SAN File System. However, you should still be able to compare the results with version 1 of the command.

1. Run the **lserver** command.

```
sfsccli lserver > lserver_v2
```

The file should contain results similar to the following:

Note: The results will not match the results you obtained when you ran this command before upgrading.

Name	State	Server Role	Filesets	Last Boot
mds2	Online	Master		5 Apr 05, 2004 5:52:45 PM
mds1	Online	Subordinate		0 Apr 05, 2004 8:18:32 PM

2. Run the **lsadmuser** command.

```
sfsccli lsadmuser > lsadmuser_v1
```

The file should contain results similar to the following:

Name	User	Role	Authorization
monitor	Monitor	Not	Current
admin	Admin	Not	Current
MDMSUID	Admin	Not	Current
operator	Operator	Not	Current
backup	Backup	Not	Current
itpc	Admin	Not	Current
superuser	Admin	Not	Current
test	Admin	Current	

3. Run the **lsautorestart** command.

```
sfsccli lsautorestart > lsautorestart_v2
```

The file should contain results similar to the following:

Name	Service	State	Probe	State	Last Probe	Probes	Highest	Retries
ivt1-mds2	Running	Live	Server	Apr 06, 2004	2:49:34 PM	7477		0
ivt1-mds1	Running	Live	Server	Apr 06, 2004	2:57:13 PM	6658		0

4. Run the **lsclient** command.

```
sfsccli lsclient > lsclient_v2
```

The file should contain results similar to the following:

Client	Session ID	State	Server	Renewals	Privilege
ivt1-c3-aix.bvnssg.net	3	Current	ivt1-mds2	54	Standard
ivt1-c4-aix.bvnssg.net	2	Current	ivt1-mds2	1545	Standard
ivt1-c2-win	1	Current	ivt1-mds2	2375	Root
ivt1-c1-win	6	Current	ivt1-mds2	10646	Root
ivt1-c3-aix.bvnssg.net	4	Current	ivt1-mds1	54	Standard

```

ivt1-c4-aix.bvnssg.net      3 Current ivt1-mds1    1548 Standard
ivt1-c2-win                 2 Current ivt1-mds1    2381 Root
ivt1-c1-win                 7 Current ivt1-mds1    9463 Root

```

- Run the **lsfileset** command (compare the results of this command with the **lscontainer** command).

```
sfscli lsfileset > lsfileset_v2
```

The file should contain results similar to the following:

```

Name      Fileset State Quota Type Quota (MB) Used (MB) Used (%) Threshold (%) Most...
=====
ROOT      Attached Soft      0      16      0      0      -
CONTAINER1 Attached Soft      0     5232      0     80      -
CONTAINER2 Attached Soft      0     3952      0     80      -
CONTAINER3 Attached Soft      0     3920      0     80      -
CONTAINER4 Attached Soft      0     3904      0     80      -

```

- Run the **lslun** command.

```
tanktool lslun > lslun_v2
```

The file should contain results similar to the following:

```

Lun ID                                     Vendor Product Size (MB) Volume State
=====
VPD83Type1=IBM 2105          20C27074 IBM 2105800 9536 vol_14 Assigned
VPD83Type1=IBM 2105          20B27074 IBM 2105800 9536 vol_13 Assigned
VPD83Type1=IBM 2105          20927074 IBM 2105800 9536 vol_09 Assigned
VPD83Type1=IBM 2105          20A27074 IBM 2105800 9536 vol_10 Assigned
VPD83Type1=IBM 2105          20827074 IBM 2105800 9536 vol_08 Assigned
VPD83Type1=IBM 2105          20727074 IBM 2105800 9536 vol_07 Assigned
VPD83Type1=IBM 2105          20627074 IBM 2105800 9536 vol_06 Assigned
VPD83Type1=IBM 2105          20527074 IBM 2105800 9536 vol_05 Assigned
VPD83Type1=IBM 2105          20427074 IBM 2105800 9536 vol_04 Assigned

```

- Run the **lspool** command.

```
tanktool lspool > lspool_v2
```

The file should contain results similar to the following:

```

Name      Type      Size (MB) Used (MB) Used (%) Threshold (%) Volumes
=====
SYSTEM    System    71520     1616      2      80      5
DEFAULT_POOL User Default 9536      16      0      80      1
TEST_POOL1 User      143040    9184     6      90      12
TEST_POOL2 User      157344    7824     4      90      12

```

- Run the **lsvol** command.

```
tanktool lsvol > lsvol_v2
```

The file should contain results similar to the following:

```

Name      State      Pool      Size (MB) Used (MB) Used (%)
=====
MASTER    Activated  SYSTEM    14304     336      2
SYSTEM1    Activated  SYSTEM    14304     336      2
SYSTEM2    Activated  SYSTEM    14304     320      2
SYSTEM3    Activated  SYSTEM    14304     320      2
SYSTEM4    Activated  SYSTEM    14304     304      2
vol_00     Activated  DEFAULT_POOL 9536      16      0
vol_01     Activated  TEST_POOL1 9536      768      8
vol_02     Activated  TEST_POOL1 9536      752      7
vol_03     Activated  TEST_POOL1 9536      736      7
vol_04     Activated  TEST_POOL1 9536      752      7

```

Performing a backup of the new master

This topic describes how to perform a backup of the new master.

After you have upgraded the initial metadata server and made it the new master metadata server, you need to make a backup of the new configuration. Then, if

you encounter a problem while upgrading any of the remaining metadata server engines, you can revert back to this image.

To backup the new master metadata server, see Chapter 6, “Backing up the SAN File System,” on page 111.

Upgrading additional metadata servers

This topic provides the procedure for upgrading all remaining metadata servers in the SAN File System cluster.

After you have upgraded the initial metadata server and made it the new master metadata server, perform these steps on *each* of the other metadata servers in the cluster.

1. Prepare the metadata server engine for upgrade so that the SAN File System cluster continues to function while you are upgrading this metadata server engine. See “Preparing an engine for upgrade.”
2. Install all software on the metadata server engine. See “Installing software on a metadata server” on page 92.
3. Restore the configuration settings. See “Restore configuration settings” on page 84.
4. Set up the metadata server engine and validate that it is installed correctly by having it rejoin the cluster and ensuring that this metadata server engine has access to LUNs and volumes. See “Setting up a metadata server engine” on page 92.
5. If you disabled the boot drive mirror, you can now re-enable the mirror. See “Re-enabling the boot drive mirror” on page 85.
6. Complete the upgrade process by reassigning filesets to this metadata server and configuring SNMP traps. See “Completing the upgrade process” on page 93.

Preparing an engine for upgrade

This topic describes how to prepare an engine to be upgraded.

These steps are the same as the steps used to prepare the first engine for upgrade.

1. Obtain software to be used on the upgrade process. See “Obtain prerequisite software” on page 30
2. Stop the engine. See “Stopping a metadata server” on page 81
3. Disable RSA II timer watchdogs. See “Disabling the RSA II watchdogs” on page 81
4. Optionally, reassign filesets to other engines. See “Reassign filesets to another metadata server in the cluster” on page 81
5. Copy configuration files to another location. See “Copy configuration files to a backup location” on page 82
6. Upgrade the system BIOS. See “Upgrading system BIOS” on page 30.
7. Upgrade the RSA II firmware. See “Upgrading RSA II firmware” on page 42.
8. Optionally, split the boot drive mirror. See “Splitting the boot drive mirror” on page 83.

Installing software on a metadata server

This topic provides an overview of the steps required to install the software on a metadata server to be upgraded.

The steps for installing the software are the same as the steps for installing software on the first metadata server engine.

1. Install SUSE LINUX Enterprise Server 8.0. See “Installing the operating system” on page 31
2. Disable the automatic starting of X-Windows (if you choose to have graphical mode as the default desktop setting during the installation of the operating system). See “Disabling the automatic starting of the X Window System” on page 34.
3. Set the time and date on the metadata server engine. See “Setting the time and date on the Metadata servers” on page 34.
4. Apply the United Linux Service Pack 3. See “Apply United Linux Service Pack 3 (SP3) updates” on page 34
5. Install all prerequisite software, such as the QLogic driver and the IBM Director agent. See “Installing prerequisite software” on page 35
6. Install the SAN File System software. See “Installing SAN File System software” on page 40
7. Restore configuration settings from your backup location. See “Restore configuration settings” on page 84

Setting up a metadata server engine

This topic describes how to set up the metadata server engine after you have installed all software.

The master metadata server is currently running version 2.1 of the SAN File System software. Therefore, you must run Version 2.1 command syntax when running any commands from the upgraded master metadata server.

1. Ensure that the engine is powered ON and that you are logged in as root.
2. Start the metadata server.

```
/usr/tank/admin/bin/setupsfs -noprompt
```
3. From the master metadata server, list all servers in the SAN File system cluster to verify that the metadata server has rejoined the cluster.

```
/usr/tank/admin/bin/sfsccli lserver
```

Note: You may have to start the upgraded metadata server. If you need to start it, you can run the **startserver** command.

```
/usr/tank/admin/bin/sfsccli startserver metadata_server_name
```

You should see the upgraded server in the list of metadata servers:

Name	State	Server Role	Containers	Last Boot
mstr_mds	Online	Master		3 Feb 13, 2004 2:46:28 PM
sub_mds1	Online	Subordinate		1 Feb 16, 2004 4:26:08 AM

If you did not move the workload from the metadata server before upgrading, the filesets managed by the upgraded metadata server should again be accessible to SAN File System client.

Completing the upgrade process

This topic provides an overview of the steps required to complete the upgrade process.

1. If you disabled the boot drive mirror, you can re-enable it. See “Re-enabling the boot drive mirror” on page 85.
2. Optionally, reassign filesets to the upgraded metadata server. See “Reassign filesets to another metadata server in the cluster” on page 88.
3. Configure SNMP traps on the upgraded metadata server. See “Configuring metadata servers for SNMP traps” on page 54.

Upgrading SAN File System on a Windows client

This topic provides the general steps for upgrading the SAN File System on a Windows client. These steps must be performed on each Windows client in the SAN File System.

1. Obtain the version 2.1 SAN File System client software. See “Obtain version 2.1 software for a Windows client” on page 47.
2. Prepare the Windows client for upgrading by stopping all applications on the client, uninstalling the current version of the SAN File System, and recording the client configuration information. See “Preparing a Windows client for upgrading” on page 94.
3. Install the version 2.1 client software. See “Installing the SAN File System software” on page 48.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on a Windows client” on page 48.
5. Set up the client to automatically start when the client is rebooted. See “Automate client restart on reboot” on page 94.

Windows client upgrade checklist

Print and use the following checklist to assist you in upgrading all of the SAN File System Windows clients.

Checklist

Windows clients			
Client host name			
Metadata server IP address			
SAN File System port			
SAN File System preferred drive letter			
Network connection type			
Critical error handling policy			
Prepare for upgrade			
Client package loaded on client			
Directory for client package			
Stop applications on client			
Stop SAN File System client			
Uninstall current SAN File System client			

Install client software			
Install client package			
Validate installation			

Preparing a Windows client for upgrading

This topic provides an overview of the tasks required to prepare a Windows client to be upgraded to version 2.1 of the SAN File System.

1. Stop all client applications that are currently running on the client.
2. Remove the current version of the SAN File System Windows client software.
3. Make a copy of the registry before upgrading the client.
4. Record the current SAN File System client configuration information.
 - SAN File System server
 - SAN File System server port
 - SAN File System preferred drive letter
 - SAN File System client name
 - SAN File System network connection type
 - SAN File System client critical error handling policy

Stop all applications on the client

This topic provides describes how to stop all applications (including SAN File System) on the Windows client.

All applications that are currently running on the SAN File System client need to be stopped. For applications other than the SAN File System, refer to the documentation that comes with the application for information about stopping it.

The SAN File System client for Windows automatically starts when you boot the client. You must modify the registry to disable the automatic restart.

1. Start the registry editor.
c:\>regedit
2. Edit the registry key
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\STFS\Start
3. Change the value of this key from 2 (the default) to 3, which stops the SAN File System client from starting when you boot the client.
4. Reboot the client

Automate client restart on reboot

This topic describes how to optionally configure the SAN File System Windows client to automatically restart when the client is rebooted:

1. Start the registry editor
c:\>regedit
2. Edit the registry key
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\STFS\Start
3. Change the value of this key from 3 to 2, which automatically starts the SAN File System client when you boot the client.
4. Reboot the client

Upgrading SAN File System on an AIX client

This topic provides the general steps for upgrading the SAN File System on an AIX client. These steps must be performed on each AIX client in the SAN File System.

1. Obtain the version 2.1 SAN File System client software. See “Obtain version 2.1 software for an AIX client” on page 49.
2. Prepare the AIX client for upgrading, which includes stopping all applications on the client and removing the current version of the SAN File System. See “Preparing an AIX client for upgrading”
3. Install the version 2.1 client software. See “Installing SAN File System on an AIX client” on page 96.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on an AIX client” on page 51.

AIX client upgrade checklist

Print and use the following checklist to assist you in upgrading all of the SAN File System AIX clients.

Checklist

AIX clients			
Client host name			
Temporary location of stclient.conf (assuming you saved stclient.conf)			
Prepare for upgrade			
Client package loaded on client			
Directory for client package			
Stop applications on client			
Stop SAN File System client			
Install client software			
Install client package			
Copy stclient.conf back to /usr/tank/client/config/			
Validate installation			

Preparing an AIX client for upgrading

This topic provides an overview of the tasks required to prepare an AIX client to be upgraded to version 2.1 of the SAN File System.

1. Stop all client applications that are currently running on the client.
2. Copy the existing client configuration file to a temporary directory on the client. The existing configuration file is /usr/tank/client/config/stclient.conf. After you have upgraded the client, you can copy the configuration file back to /usr/tank/client/config.

Note: This file will exist only if you saved the configuration the last time you ran the **setupstclient** command.

Stop all applications on the client

This topic provides describes how to stop all applications (including SAN File System) on the AIX client.

All applications that are currently running on the SAN File System client need to be stopped. For applications other than the SAN File System, refer to the documentation that comes with the application for information about stopping it.

Perform these steps to stop the SAN File System client on AIX:

1. Telnet to the client from the master console.
`telnet client_name`
2. Log in to the client.
3. Determine the current mount point for the client.
`mount`
4. Verify that the SAN File System is not currently in use
`fuser -u mount_point`

Note: You cannot unmount the client if the SAN File System is in use.

5. Stop and unmount the SAN File System client for AIX.
`/usr/tank/client/bin/rmstclient`

This command unmounts the global namespace, stops the client, and unloads the kernel module. It uses the configuration information stored in `/usr/tank/client/config/stclient.conf`.

6. Verify that the client is gone. The mount point for the SAN File System should no longer be displayed.
`mount`

Installing SAN File System on an AIX client

This topic describes how to install version 2.1 of the SAN File System on an AIX client

- You must have root privileges to install the client for AIX.
 - For AIX 5.1, the `bos.mp` (multiprocessor) or `bos.up` (uniprocessor) packages must be at least 5.1.0.58 or higher. AIX 5.1 32-bit high availability cluster multi-processing (HACMP) environments are supported at the specified maintenance level.
 - For AIX 5.2, the `bos.mp`, `bos.up`, or `bos.mp64` must be at least 5.2.0.18 or later.
1. Enable asynchronous input/output on the AIX client, which is required for the SAN File System.
 - a. Log on to the client as root.
 - b. Start **smit**.
 - c. Select **Devices**.
 - d. Select **Asynchronous I/O**.
 - e. Select **Asynchronous I/O (Legacy)**. You will see this choice only if you are enabling asynchronous I/O on AIX 5.2. If you are using AIX 5.1, skip this step.
 - f. Select **Change/Show Characteristics of Asynchronous I/O**
 - g. Use the tab key to set the STATE to be configured at system restart to **available**.
 - h. Press Enter.

- i. Exit smit.
- j. Run `cfgmgr` to apply the changes.
2. Navigate to the directory where the client installation package is located.
3. Use `smit` or `installp` to install the package. For example:


```
installp -ac -d . client_package_name
```
4. Configure and start the client. If you are using the client configuration as it existed before the upgrade, follow these steps:
 - a. Copy `stclient.conf` from the temporary directory back to `/usr/tank/client/config/`.
 - b. Run the `setup` command


```
/usr/tank/client/bin/setupstclient
```
5. If you do not have `stclient.conf` or if you want to make changes to the client configuration, run the `setup` command with the `-prompt` parameter.


```
/usr/tank/client/bin/setupstclient -prompt
```

You will be prompted to enter values for the client configuration. In most cases you can accept the defaults. If you change these values, make sure that you type the new values correctly.

Verifying the installation

This task describes ways to verify various parts of the installation.

1. Clients: use `ping` to validate connectivity between clients.
2. Metadata servers:
 - a. Validating connectivity between the metadata servers and the master console:
 - 1) Use `ping` to validate connectivity.
 - 2) Use SSH to validate connectivity.
 - 3) Use a browser from the master console to connect to the SAN File System console.
 - b. Show engines in the cluster.


```
/usr/tank/admin/bin/sfsccli lsserver
```
 - c. Show clients connected to the server.


```
/usr/tank/admin/bin/sfsccli lsclient
```
 - d. Verify that the disk devices that are intended to be used by the data and metadata storage pools are visible and accessible on each node.
 - e. Validating security: you can validate the privileged clients with `lsserverconfig`.
 - f. Validating policies: you can view the policy rules.
 - g. Validating alerts and SNMP traps: you can validate the alerts set by using the `lsserverconfig` parameter.

Upgrading the cluster

This topic describes how to upgrade the cluster.

Before you upgrade the cluster, you must ensure that all metadata servers in the SAN File System cluster and all SAN File System clients are upgraded to version 2.1.

1. From the master metadata server, activate the version 2.1 upgrade to enable new version 2.1 features.
`/usr/tank/admin/bin/sfscli upgradeclass`
2. When prompted to confirm that you want to upgrade, respond with **Y**
3. Verify that the version 2.1 level of code is now the committed software version.
`/usr/tank/admin/bin/sfscli statclass`

The entries for Software Version and Committed Software Version should be the same and should match the version name of the SAN File System Server package that was installed. Run `rpm -q sfs.server.linux` and verify the versions match.

```
Name                sanfs
ID                  21443
State               Online
Target State        Online
Last State Change   Feb 16, 2004 6:15:47 AM
Last Target State Change -
Servers              2
Active Servers      1
Software Version    2.1.0.65
Committed Software Version 2.1.0.65
Last Software Commit Feb 13, 2004 2:46:27 PM
Software Commit Status Not In Progress
Installation Date    Feb 13, 2004 2:46:27 PM
```

The results of `rpm -q sfs.server.linux` should be similar to the following:

```
# rpm -q sfs.server.linux
sfs.server.linux-2.1.0-65
```

Revising scripts and policies

This topic provides considerations for updating the zoning of the metadata server zones.

If you created scripts based on the version 1 administrative command-line syntax, you may need to modify those scripts to match the command syntax for version 2. These changes include modifications to the command names as well as modifications to the log messages.

In addition, you may also need to modify your policies if you defined placement rules based on containers.

The following terminology changes have been implemented for version 2:

	Version 1.1	Version 2.1
Administrative CLI	tanktool	sfscli
Container commands have been renamed to fileset commands	mkcontainer	mkfileset
	lscontainer	lsfileset
	chcontainer	chfileset
	rmcontainer	rmfileset
	attachcontainer	attachfileset
	detachcontainer	detachfileset
	setcontainerserver	setfilesetserver
Setting up SAN File System	setupTank	setupfs

Modifying metadata server zoning and LUN masking

This topic provides considerations for updating the zoning of the metadata server zones and LUN masking.

In version 1 of the SAN File System, it was recommended that you implement a zone containing the metadata servers in the cluster, the volumes (LUNs) comprising the system storage pool, and volumes comprising the user storage pools. In version 2 of the SAN File System, metadata servers no longer need access to the volumes comprising the user storage pools. It is recommended, but not required, that you modify the metadata server zones to remove all user data LUNs.

If you are using LUN masking, you want to unassign the volumes comprising the user storage pools from the metadata servers.

Attention: You must not unassign or zone off a system storage pool LUN from the metadata server or you will lose your SAN File System cluster.

Refer to the documentation for your SAN switches and storage subsystems for information about zoning.

Chapter 5. Upgrading the SAN File System from Version 2.1

This topic provides an overview of the process for upgrading version 2.1 of the SAN File System to a later release.

Before you begin upgrading the SAN File System, make sure that you have access to the SAN File System CD-ROM.

Use the following checklist to upgrade from version 2.1:

Steps		For more information...
1	Create a backup of your existing system.	Chapter 6, "Backing up the SAN File System," on page 111
2	Upgrade the master console.	<i>Master Console User's Guide</i>
3	Upgrade the package repository on each metadata server engine in the cluster.	"Upgrading the package repository" on page 103
4	Upgrade each subordinate metadata server engine.	"Upgrading the first metadata server engine" on page 80
	a Prepare the engine for upgrade.	"Preparing the metadata server for upgrade" on page 104
	b Upgrade the administrative server package.	"Upgrading the administrative server package" on page 104
	c Upgrade the metadata server package.	"Upgrading the metadata server package" on page 104
	d Restart the metadata server engine.	"Starting the metadata server engine" on page 105
5	Upgrade the master metadata server engine.	"Upgrading metadata server engines" on page 103
	a Upgrade the administrative server package.	"Upgrading the administrative server package" on page 104
	b Upgrade the metadata server package.	"Upgrading the metadata server package" on page 104
	c Restart the metadata server engine.	"Starting the metadata server engine" on page 105
6	Upgrade clients.	
	a Upgrade Windows clients.	"Upgrading SAN File System on a Windows client" on page 105
	1 Obtain Windows client software.	"Obtain version 2.1 software for a Windows client" on page 47
	2 Prepare the client for upgrade.	"Preparing a Windows client for upgrading" on page 94
	a Stop all client applications.	"Stop all applications on the client" on page 94
	b Uninstall the Windows client	"Uninstalling the SAN File System software from a Windows client" on page 69
	3 Install the software.	"Installing the SAN File System software" on page 48

Steps		For more information...
4	Validate client installation.	"Validating the installation of SAN File System on a Windows client" on page 48
5	Set up the client to automatically start when the client is rebooted.	"Automate client restart on reboot" on page 94
b	Upgrade AIX clients.	"Upgrading SAN File System on an AIX client" on page 105
1	Obtain AIX client software.	"Obtain version 2.1 software for an AIX client" on page 49
2	Prepare the client for upgrade.	"Preparing an AIX client for upgrading" on page 95
a	Stop all client applications.	"Stop all applications on the client" on page 96
3	Install the software.	"Installing the SAN File System software" on page 50
4	Validate client installation.	"Validating the installation of SAN File System on an AIX client" on page 51
c	Upgrade Linux clients.	"Upgrading SAN File System on a Linux client" on page 106
1	Obtain Linux client software.	"Obtain version 2.1 software for a Linux client" on page 51
2	Prepare the client for upgrade.	"Preparing a Linux client for upgrading" on page 106
a	Stop all client applications.	"Stop all applications on the client" on page 106
3	Install the software.	"Upgrading the SAN File System software" on page 107
4	Validate client installation.	"Validating the installation of SAN File System on a Linux client" on page 52
d	Upgrade Solaris clients.	"Upgrading SAN File System on a Solaris client" on page 107
1	Obtain Solaris client software.	"Obtain version 2.1 software for a Solaris client" on page 53
2	Prepare the client for upgrade.	"Preparing a Solaris client for upgrading" on page 108
a	Stop all client applications.	"Stop all applications on the client" on page 108
3	Install the software.	"Installing the SAN File System software" on page 53
4	Validate client installation.	"Validating the installation of SAN File System on a Solaris client" on page 54
7	Upgrade the cluster.	"Committing the upgrade" on page 109
8	Back up the complete system.	Chapter 6, "Backing up the SAN File System," on page 111

Upgrading the package repository

This topic describes how to upgrade the SAN File System package repository on a metadata server engine.

You must be logged in with root privileges to upgrade the package repository.

The SAN File System package repository holds all the packages needed to install the various SAN File System software components, including the metadata server, the administrative server, and all clients. By default, these packages are installed in `/usr/tank/packages`.

You need to perform this procedure on each metadata server engine in the cluster.

Note:

- A single up-to-date package repository can be used to serve packages to the entire SAN File System. However, to avoid the accidental installation of down-level packages and for high availability, it is recommended that all copies of the package repository be kept up to date.
- If the name of the updated package has not changed since the previous version, the package is overwritten. To keep backup copies of old packages, you should copy them from `/usr/tank/packages` to some other location.

1. Make sure the CD-ROM drive is mounted and insert the SAN File System CD-ROM into the CD-ROM drive.
2. If you have all engines attached to a single KVM, switch the monitor to this engine. Otherwise, establish an SSH session from the master console to the engine.
3. Determine the version of the currently installed package repository.

```
rpm -qa |grep sfs
```

4. If the new package repository name is different than the existing package repository, remove the existing package repository.

```
rpm -e existing_package_repository_name
```

5. Install the new package repository.

```
rpm -Uvh new_package_repository_name
```

For example:

```
rpm -Uvh sfs-package
```

Note: To review the contents of the package before installing it, use the following command:

```
rpm -qpl package_repository_name
```

Continue with the next metadata server in the cluster until you have upgrade the package repository on all metadata server engines.

Upgrading metadata server engines

This topic provides the general steps for upgrading metadata server engines from version 2.1.

You will need to perform these procedures for all metadata servers in the cluster.

Important: Start by upgrading all subordinate metadata servers. Then upgrade the master metadata server.

1. Prepare the server to be upgraded by stopping the server.
2. Upgrade the administrative server on that engine.
3. Upgrade the metadata server on that engine.
4. Verify the upgrade process and restart the engine.

Complete this procedure for each of the subordinate metadata servers in the cluster. Then upgrade the master metadata server.

Preparing the metadata server for upgrade

This topic describes how to check the current state of the cluster and stop the metadata server that you are upgrading.

1. From the master metadata server engine, change directories to the binaries directory.

```
cd /usr/tank/admin/bin
```
2. Check the current state of the cluster.

```
./sfsccli statcluster
```
3. Start the Administrative command-line interface in interactive mode.

```
./sfsccli
```
4. Stop the metadata server that you are going to upgrade.

```
stopserver server_name
```
5. When prompted to confirm that you want to stop this server, respond with yes.
6. Verify that the server has stopped.

```
lsserver server_name
```

The selected server should be listed as "Not Running."

Upgrading the administrative server package

This topic describes how to upgrade the administrative server package from release 2.1.

1. If you are accessing the metadata servers from a single KVM, switch to the metadata server engine you are upgrading. Otherwise, establish an SSH session between the master console and the engine you are upgrading.
2. Upgrade the administrative server package.

```
rpm -Uvh administrative_package_name
```

For example:

```
rpm -Uvh /usr/tank/packages/sfs.admin.linux
```

Upgrading the metadata server package

This topic describes how to upgrade the metadata server package from release 2.1.

This procedure assumes that you have already switched to the engine you are upgrading (if you are accessing the metadata servers from a single KVM) or that you have established an SSH session between the master console and the engine that you are upgrading. In addition, it assumes that the metadata server is already stopped.

1. Stop the Cimom agent.

```
/usr/tank/admin/bin/stopCimom
```

2. Upgrade the metadata server package.
`rpm -Uvh metadata_server_package_name`
For example:
`rpm -Uvh /usr/tank/packages/sfs-server-linux`
3. Restart the Cimom agent.
`/usr/tank/admin/bin/startCimom`

Starting the metadata server engine

This topic describes how to verify the upgrade process and restart the metadata server engine.

This procedure assumes that you have already switched to the engine you are upgrading (if you are accessing the metadata servers from a single KVM) or that you have established an SSH session between the master console and the engine that you are upgrading.

1. From the upgraded engine, verify that the package was upgraded successfully.
`rpm -qa|grep sfs`
2. From the sfscli prompt on the master metadata server, restart the upgraded metadata server.
`startserver upgrade_metadata_server_name`
3. When prompted to confirm that you want to start this engine, respond with yes.

Upgrading SAN File System on a Windows client

This topic provides the general steps for upgrading the SAN File System on a Windows client. Perform these steps on each Windows client in the SAN File System.

1. Obtain the latest version of the SAN File System client software. See “Obtain version 2.1 software for a Windows client” on page 47.
2. Prepare the Windows client for upgrading by stopping all applications on the client, removing the current version of the SAN File System, and recording the client configuration information. See “Preparing a Windows client for upgrading” on page 94.
3. Install the latest version of the client software. See “Installing the SAN File System software” on page 48.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on a Windows client” on page 48.

Upgrading SAN File System on an AIX client

This topic provides the general steps for upgrading the SAN File System on an AIX client. Perform these steps on each AIX client in the SAN File System.

1. Obtain the latest version of the SAN File System client software. See “Obtain version 2.1 software for an AIX client” on page 49.
2. Prepare the AIX client for upgrading by stopping all applications on the client. See “Preparing an AIX client for upgrading” on page 95.
3. Install the latest version of the client software. See “Installing the SAN File System software” on page 50.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on an AIX client” on page 51.

Upgrading SAN File System on a Linux client

This topic provides the general steps for upgrading the SAN File System on a Linux client. Perform these steps on each Linux client in the SAN File System.

1. Obtain the latest version of the SAN File System client software. See “Obtain version 2.1 software for a Linux client” on page 51.
2. Prepare the Linux client for upgrading by stopping all applications on the client. See “Preparing a Linux client for upgrading.”
3. Install the latest version of the client software. See “Upgrading the SAN File System software” on page 107.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on a Linux client” on page 52.

Linux client upgrade checklist

Print and use the following checklist to assist you in upgrading all of the SAN File System Linux clients.

Checklist

Linux clients			
Client host name/IP address			
Temporary location of stclient.conf			
Prepare for upgrade			
Client package loaded on client			
Directory for client package			
Stop applications on client			
Stop SAN File System client			
Install client software			
Upgrade client package			
Validate installation			

Preparing a Linux client for upgrading

This topic provides an overview of the tasks required to prepare a Linux client to be upgraded to version 2.1 of the SAN File System.

1. Stop all client applications that are currently running on the client. See “Stop all applications on the client.”
2. Copy the existing client configuration file to a temporary directory on the client for backup. The existing configuration file is
/usr/tank/client/config/stclient.conf.

Stop all applications on the client

This topic provides describes how to stop all applications (including SAN File System) on the Linux client.

All applications that are currently running on the SAN File System client need to be stopped. For applications other than the SAN File System, refer to the documentation that comes with the application for information about stopping it.

Perform these steps to stop the SAN File System client for Linux.

1. Verify that the SAN File System is not currently in use.

```
fuser -u mount_point
```

You cannot unmount the client if the SAN File System is in use.

2. Stop and unmount the SAN File System client for Linux.

```
/usr/tank/client/bin/rmstclient
```

This command unmounts the global namespace, stops the client, and unloads the kernel module. It will use the configuration information stored in `/usr/tank/client/config/stclient.conf`.

Upgrading the SAN File System software

This topic describes how to install the SAN File System on a Linux client

1. Navigate to the directory where the client installation package is located.

2. Install the client package.

```
rpm -U sfs.client.linux_RHEL3_9-build_level.i386.rpm
```

3. Make sure that the master metadata server is running.

4. Configure and start the client. Run the setup command.

```
/usr/tank/client/bin/setupstclient -prompt
```

You will be prompted to enter values for the client configuration.

- SAN File System server name (no default)
- SAN File System server port (default is 1700)
- SAN File System mount point (no default)
- SAN File System client name (default is the short version of the hostname)
- SAN File System network connection type (default is TCP)
- SAN File System client critical error handling policy (default is log)

In most cases you can accept the defaults.

Tip: If you have installed SDD on the client, you should use the following device pattern when prompted for storage devices.

```
pat=/dev/vpath*[a-z]
```

Upgrading SAN File System on a Solaris client

This topic provides the general steps for upgrading the SAN File System on a Solaris client. Perform these steps on each Solaris client in the SAN File System.

1. Obtain the latest version of the SAN File System client software. See “Obtain version 2.1 software for a Solaris client” on page 53.
2. Prepare the Solaris client for upgrading by stopping all applications on the client. See “Preparing a Solaris client for upgrading” on page 108.
3. Install the latest version of the client software. See “Installing the SAN File System software” on page 53.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on a Solaris client” on page 54.

Solaris client upgrade checklist

Print and use the following checklist to assist you in upgrading all of the SAN File System Solaris clients.

Checklist

Solaris clients			
Client host name/IP address			
Temporary location of stclient.conf			
Prepare for upgrade			
Client package loaded on client			
Directory for client package			
Stop applications on client			
Stop SAN File System client			
Install client software			
Install client package			
Validate installation			

Preparing a Solaris client for upgrading

This topic provides an overview of the tasks required to prepare a Solaris client to be upgraded to version 2.1 of the SAN File System.

1. Stop all client applications that are currently running on the client. See “Stop all applications on the client.”
2. Optionally, uninstall the current version of the SAN File System Solaris client software. See “Uninstalling the SAN File System software from a Solaris client” on page 70.
3. Copy the existing client configuration file to a temporary directory on the client for backup. After you have upgraded the client, you can copy the configuration file back to /usr/tank/client/config.

Stop all applications on the client

This topic provides describes how to stop all applications (including SAN File System) on the Solaris client.

All applications that are currently running on the SAN File System client need to be stopped. For applications other than the SAN File System, refer to the documentation that comes with the application for information about stopping it.

Perform these steps to stop the SAN File System client for Solaris.

1. Verify that the SAN File System is not currently in use.

```
fuser -u mount_point
```

Note: You cannot unmount the client if the SAN File System is in use.

2. Stop and unmount the SAN File System client for Solaris.

```
/usr/tank/client/bin/rmstclient
```

This command unmounts the global namespace, stops the client, and unloads the kernel module. It uses the configuration information stored in `/usr/tank/client/config/stclient.conf`.

Committing the upgrade

This topic describes how to upgrade the cluster.

Before you upgrade the cluster, you must ensure that all metadata servers in the SAN File System cluster and all SAN File System clients are upgraded.

1. Activate the upgrade to enable new features.
`/usr/tank/admin/bin/sfsccli upgradecoluster`
2. When prompted to confirm that you want to upgrade, respond with yes.
3. Verify that the upgraded version is now the committed software version.
`/usr/tank/admin/bin/sfsccli statcluster`

The entries for Software Version and Committed Software Version should match.

Chapter 6. Backing up the SAN File System

This section directs you to the appropriate sections for backing up the parts of the SAN File System.

Backing up using the LUN method

Go to “Backing up using the LUN method” for information about how to perform a backup using the LUN-based approach.

Backing up using the file-based (API) method

Go to “Backing up using the file-based (API) method” on page 112 for information about how to perform a backup using the file-based approach.

Saving a copy of metadata server and client configuration

Go to “Saving additional SAN File System configuration files” on page 114 for creating FlashCopy images of selected filesets.

Backing up filesets

Go to “Saving a FlashCopy image of a fileset” on page 115 for creating FlashCopy images of selected filesets.

Managing backups

SAN File System supports the use of backup tools that are already present in your environment. For example, if your enterprise currently uses a storage management product such as Tivoli® Storage Manager (TSM), you can use the functions and features of that product to back up and restore files that reside in the SAN File System global namespace.

For backing up in a normal, available environment, you can use the FlashCopy image feature of SAN File System.

To prepare for disaster recovery in situations where SAN File System becomes unavailable, you can perform LUN-based backups using the instant copy features that exist in the storage subsystems that SAN File System supports. If your SAN storage subsystems do not offer copy services, you must back up for disaster recovery using the API method.

Backing up using the LUN method

This topic describes how to perform SAN File System backup operations using the LUN method. LUN backup requires that all transactions are stopped during the process.

The LUN method of backup is only available to SANs comprised of storage subsystems with built-in copy services. SANs without such service must use the file-based (API) method of backup.

Because the LUN method deals with data at the byte level, it is an all-or-nothing approach for backing up and restoring your entire SAN File System. In particular, it provides no ability to restore individual files (because it has no concept of files); you have to save and restore all the data — metadata and file data — or none of it. Restoring a previously saved FlashCopy image is the best method for recovering some subset of SAN File System data. Therefore, the LUN method is best employed as part of a disaster recovery situation.

1. Stop or pause all SAN File System client applications. Because this task is application-specific, refer to the application documentation for details on performing this step.
2. The metadata server and all clients must complete all active transactions and flush their data to disk. From the master metadata server, quiesce the SAN File System metadata servers.

```
/usr/tank/admin/bin/sfsccli quiescecluster -state full
```

This procedure will also lock out any subsequent new I/O from the clients or metadata server.

3. Initiate the storage subsystem copy service using the procedure defined in its accompanying documentation.
4. After the storage subsystem copy is complete, re-enable the SAN File System metadata servers.

```
/usr/tank/admin/bin/sfsccli resumecluster
```

This procedure will also lock out any subsequent new I/O from the clients or metadata server.

5. Restart the client applications using the specific procedures for those applications.

For additional information about restore procedures, including commands, refer to the *IBM TotalStorage SAN File System Maintenance and Problem Determination Guide*, on the publications CD that came with your metadata servers.

Backing up using the file-based (API) method

This topic describes how to perform SAN File System backup operations using two variations of the file-based (API) method.

The file method of backup is used for SANs consisting of storage subsystems that do not offer built-in copy services. SANs that do offer copy services can use the LUN method of backup.

You have two possible options when using the file method of backup. Which method you choose depends on the characteristics of the backup application in your existing environment.

- If your existing backup application allows you to selectively choose subdirectory branches for backup, and allows you to restore files to the directory two levels above their original location, follow this optimized procedure for SAN File System file-based backup:

1. Stop or pause all SAN File System client applications. Because this task is application-specific, refer to the application documentation for details on performing this step.
2. Create FlashCopy images of each fileset.

```
/usr/tank/admin/bin/sfsccli mkimage -fileset fileset_name  
-dir directory_name Flashcopy_image_name
```

Tip: Consider using the SAN File System console. It allows you to create multiple FlashCopy images quickly and easily. See “Saving a FlashCopy image of a fileset” on page 115.

3. From the master metadata server engine, save the most recent metadata to accompany the FlashCopy images.

```
/usr/tank/admin/sfsccli mkdrfile most_recent_metadata_file_name
```

A text file is created in /usr/tank/server/DR. Copy the resulting file onto your backup medium (usually tape).

4. On each metadata server backup additional operating system and SAN File System administration configuration files.

```
/usr/tank/admin/bin/setupsfs -backup
```

A tar file called *DRfiles-metadata_server_name-timestamp.tar.gz* is created in /usr/tank/server/DR/ (by default). Copy the resulting tar file onto your backup medium (usually tape).

5. Restart the client applications using the specific procedures for those applications.
6. Use the backup application to back up all *fileset_name/directory_name* subdirectories and their contents, along with all of the *DRfiles-metadata_server_name-timestamp.tar.gz* files to your backup medium (usually tape).

Attention: Backup Windows filesets only from a Windows client; backup AIX filesets only from an AIX client.

- If your existing backup application does not provide the features required for the enhanced method, then follow this procedure for SAN File System file-based backup:

Note: The enhanced method allows you to re-enable client applications more quickly. Creating FlashCopy images is much quicker than backing up the real files to tape.

1. Stop or pause all SAN File System client applications. Because this task is application-specific, refer to the application documentation for details on performing this step.

2. From the master metadata server engine, save the most recent metadata.

```
/usr/tank/admin/bin/sfsccli mkdrfile most_recent_metadata_file_name
```

A text file is created in /usr/tank/server/DR/. Copy that file onto your backup medium (usually tape).

3. Use the backup application to backup all *fileset_name/directory_name* subdirectories and their contents, to your backup medium (usually tape). If possible, exclude all .flashcopy subdirectories and their contents because they will not be of any use during a subsequent restore operation.

4. On each metadata server backup additional operating system and SAN File System administration configuration files.

```
/usr/tank/admin/bin/setupsfs -backup
```

A tar file called *DRfiles-metadata_server_name-timestamp.tar.gz* is created in /usr/tank/server/DR/ (by default). Copy the resulting tar file onto your backup medium (usually tape).

5. Restart the client applications using the specific procedures for those applications.

For additional information about restore procedures, including commands, refer to the *IBM TotalStorage SAN File System Maintenance and Problem Determination Guide*, provided on the publications CD that came with your metadata servers.

Saving additional SAN File System configuration files

This topic provides an overview of additional SAN File System configuration files that should be saved.

In addition to the files that you back up, consider saving the output from the One Button Data Collector and the Target Machine Validation Tool.

Use the One Button Data Collector (OBDC) to save copies of server and client information for future reference and diagnostic purposes. See “One-button data collection” for instructions.

Use the Target Machine Validation Tool (TMVT) to save information about the hardware and software configuration of each metadata server engine. For example,
`/usr/tank/server/bin/tmvt -r tmvt_output_file`

For more information about tmvt, see “tmvt” on page 120.

One-button data collection

This topic describes how to use the SAN File System script for one-button data collection for servers and for clients.

1. The one-button data collection utility is designed to gather information of interest for first-failure data capture and analysis. It gathers diagnostic data of value in the initial investigation of reported problems. You should keep in mind that the amount of data gathered can be significant, so you might want to mount a file system over the output directory.
2. Each system provides unique log information which must be collected individually. The script must be executed on each server and each client. The script places a set of files into the directory indicated by the tool.
3. For metadata servers, do the following:
 - a. Change to the directory where the data collector script is located as follows:
`cd /usr/tank/server/bin`
 - b. Follow the menu options to collect the data. At the prompt enter:
`./obdc`
4. For UNIX-based clients, do the following:
 - a. Change to the directory where the data collector script is located as follows:
`cd /usr/tank/client/bin`
 - b. Follow the menu options to collect the data. At the prompt enter:
`./obdc`
5. For Windows clients, do the following:

Note: Output is stored in C:\Program Files\IBM\Storage Tank\pmf*system_name*\Administrator.

- a. At the C:\WINNT> prompt, change to the directory where the data collector script is located as follows:
`cd \Program Files\IBM\Storage Tank\Client\bin`
- b. Follow the menu options to collect the data. At the C:\WINNT> prompt enter:
`obdc`

Saving a FlashCopy image of a fileset

FlashCopy images are saved on a per-fileset basis.

- Using the GUI, select **Maintain System** in the My Work pane, then select **Create FlashCopy Images**.
- Click **Next**.
- Under Select Containers, select the filesets for which you want a FlashCopy Image.
- Click **Next**.
- Under Set Properties, accept the defaults and then click **Next**.
- Verify your settings, then click **Finish**.
- To list the FlashCopy Images, select **Maintain System** then select **FlashCopy Images**. The default FlashCopy Image name (Image-1) should be included in the list.
- Change ownership and permissions on the fileset's .flashcopy directory to navigate it. The directory contains an entry for each FlashCopy Image name.
- Change the directory representing the image name to view the files as their images were created.

Note: Attempting to write to the file causes an error stating that this area is read only. This applies to the .flashcopy directory and those directories and files below the .flashcopy directory.

Chapter 7. SAN File System installation commands

This topic provides an overview of the SAN File System installation commands.

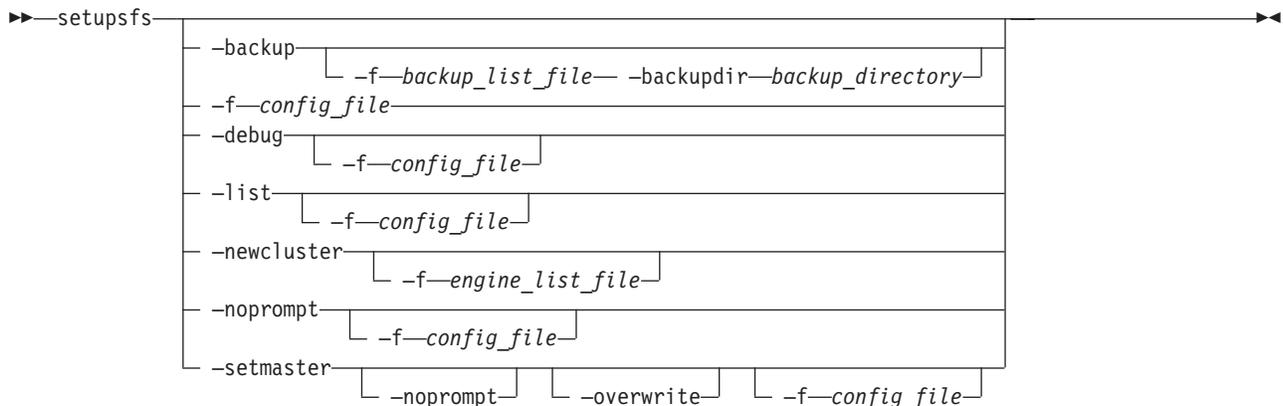
The SAN File System installation commands include:

- `Setupsfs`. Used to configure and start a SAN File System metadata server.
- `Setupstclient`. Used to configure and start a SAN File System client.
- `Tmvt`. Used to verify that the hardware and software prerequisites for the SAN File System are installed.

In addition, information about the administrative commands and all client commands are provided in the *SAN File System Administrator's Guide*.

setupsfs

Configures and starts a SAN File System metadata server



Parameters

-backup *backup_file*

Creates a backup (tar archive) containing all metadata server configuration information. The list of files to be backed up are specified in `/usr/tank/admin/config/backup.list`. You can specify an alternate backup file list using the `-f` parameter.

By default the archive is stored as `/usr/tank/server/DR/DRfiles-hostname-date.tar.gz`

To restore a metadata server from a backup that was created using the `-backup` parameter, extract the tar archive in the root directory. For example:

```
tar -xzf /usr/tank/server/DR/DRfiles-hostname-date.tar.gz
```

Then, run **setupsfs** without parameters (or with only `-noprompt` parameter) to restart the metadata server.

Attention: Do not use the `-setmaster` parameter when restoring from a backup.

- backupdir** *directory_path*
Specifies an alternate directory, in which to create the archive.
- debug**
Provides extra parameters with defaults. The most common use of this parameter is when you are using ActiveDirectory as your LDAP server.
- f** *file*
Specifies the name of a file that contains configuration information, list of engines, or tar archive, depending on the parameters specified before the **-f** parameter.
- list**
List all parameters and the corresponding values found in the configuration file.
- noprompt**
Runs the **setupsfs** command without prompting you for information. The configuration file is expected to exist. If the configuration file does not exist or if a required value is missing or invalid, **setupsfs** exits with an error.
- newcluster** *engine_list*
Adds one or more subordinate metadata servers to the cluster. Use this parameter only on the master metadata server.

The *engine_list* identifies the subordinate nodes to be added. All subordinate metadata servers must be running at the time that **setupsfs -newcluster** is started.
- overwrite**
Initialize the master metadata server and system disks, regardless of whether they already contain cluster information.

Attention: Using this parameter will destroy all data stored in the system storage pool.
- setmaster**
Configures the metadata server as the master metadata server. If the given master disk already contains cluster information, this information is used when starting the metadata server.

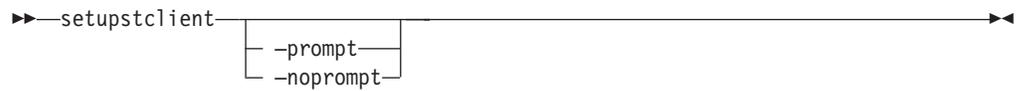
Description

The **setupsfs** command is used to configure, start, or restart a metadata server. It can also be used to display the current configuration values and back up the configuration in case the server needs to be restored at a later date.

If you do not specify any parameters, **setupsfs** prompts you for the configuration values required to set up a subordinate metadata server. The **setupsfs** command maintains the values in a configuration file in parameter=value format. The configuration file is `/usr/tank/admin/config/tank.properties`. If a configuration file already exists, its values will be presented as the suggested defaults when the prompt is displayed. Otherwise, the manufacturing defaults are presented.

setupstclient

Configures and starts SAN File System clients.



Parameters

-prompt

Forces the **setupstclient** command to prompt for all configuration values.

-noprompt

Runs silently, using parameters from the configuration file. If the configuration file does not exist or if a required parameter is not available or invalid, the command exits with an error.

Prerequisites

You must have root privileges to use this command.

Description

This command configures and starts, or restarts a SAN File System client.

If you do not specify a parameter, this command runs silently using values from the configuration file as defaults. It only prompts for any required information, if a configuration file does not exist or if a value in an existing configuration file is not valid.

This command maintains any values given by the user in a configuration file in `parameter=value` format. The default configuration file is `/usr/tank/client/config/stclient.conf`.

Specify the `-prompt` parameter to force the command to prompt for all configuration values. In this case, if a configuration file exists, the command presents the value from the configuration file as the suggested default when the command displays a prompt. If a configuration file does not exist, the command presents the manufacturing default as the suggested default.

If you specify the `-noprompt` parameter, the command expects the configuration file to exist. If the file does not contain valid values, the command exits with an error.

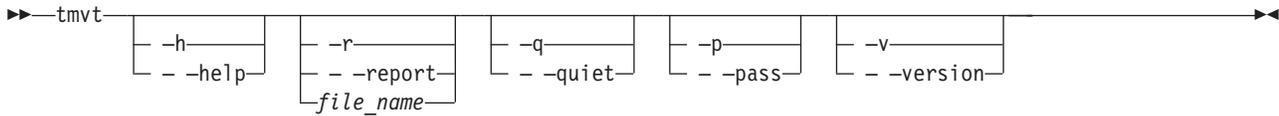
Example

Setup a client The following example configures and starts SAN File System clients:

```
setupstclient
```

tmvt

Validates that the existing hardware and software on a metadata server meet all installation requirements for the SAN File System.



Parameters

-? | -h | --help

Displays a detailed description of this command, including syntax, parameter descriptions, and examples. If you specify a help option, all other command options are ignored.

-r | --report *file_name*

The file name into which the report is written. The default is to write the report to the standard output stream, stdout.

-q | --quiet

Run in quiet mode. Does not display any output.

-p | --pass

Force the return code from tmvt processing to be zero (hardware and software passes).

-v | --version

Output the version number of the **tmvt** command and the product release that it is validating.

Description

Note: Exceptions are always written to stderr and tmvt analysis is written (by default) to stdout. If you specify a report file using the -report parameter, tmvt analysis is written to that file. Exceptions are written to the file and to stdout.

Example

The following example validates the hardware and software for a metadata server:

```
tmvt
```

The following example shows a sample of a partial listing from the file.

```
Hardware Components (15)
```

Item Name	Current	Recipe
Passed Hardware Component Checks (15)		
Memory (Megabytes)	4040	4000
Disk space in /var (Megabytes)	16472	4096
TCP/IP	enabled	enabled
Ethernet controller	Intel 82546EB	Intel 82546EB
Machine Type/Model	867061X	867061X
Machine BIOS Level	1.1.5	1.1.5
Machine BIOS Build	GEE148AUS	GEE152AUS
FC HBA Manufacturer	QLogic	QLogic
FC HBA Model	QLA2342	QLA2342
FC HBA BIOS/Firmware Version	3.02.16	3.02.16
FC HBA Driver Version	6.06.60	6.06.60

Remote Supervisor Adapter 2	present	present
Remote Supervisor Adapter 2 Firmware	GEB833A	GEB833A
Remote Supervisor Adapter 2 Driver	1.08	1.08
RS485 Service Processor	enabled	enabled

Software Components (794)

Item Name	Current	Recipe
Correct Software Packages (794)		
zsh	4.0.6-30	4.0.6-30
zoo	2.10-602	2.10-602
zlib-devel	1.1.4-51	1.1.4-51
zlib	1.1.4-51	1.1.4-51
zip	2.3-490	2.3-490
zebra	0.93b-74	0.93b-74
ypserv	2.9.91-27	2.9.91-27
ypbind	1.12-55	1.12-55
yp-tools	2.7-60	2.7-60

Appendix A. Accessibility

This topic provides information about the accessibility features of SAN File System and its accompanying documentation.

Accessibility features help a user who has a physical disability, such as restricted mobility or limited vision, to use software products successfully.

Features

These are the major accessibility features in SAN File System:

- You can use screen-reader software and a digital speech synthesizer to hear what is displayed on the screen.

Note: The SAN File System Information Center and its related publications are accessibility-enabled for the IBM Home Page Reader.

- You can operate all features using the keyboard instead of the mouse.

Navigating by keyboard

You can use keys or key combinations to perform operations and initiate many menu actions that can also be done with a mouse. You can navigate the SAN File System console and help system from the keyboard by using the following key combinations:

- To traverse to the next link, button or topic, press Tab inside a frame (page).
- To expand or collapse a tree node, press Right Arrow or Left Arrow, respectively.
- To move to the next topic node, press Down Arrow or Tab.
- To move to the previous topic node, press Up Arrow or Shift+Tab.
- To scroll all the way up or down, press Home or End, respectively.
- To go back, press Alt+Left Arrow
- To go forward, press Alt+Right Arrow.
- To go to the next frame, press Ctrl+Tab. There are quite a number of frames in the help system.
- To move to the previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.

Appendix B. Getting help, service, and information

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

IBM maintains pages on the World Wide Web where you can get information about IBM products and services and find the latest technical information.

Table 2 lists some of these pages.

Table 2. IBM Web sites for help, services, and information

www.ibm.com/	Main IBM home page
www.ibm.com/storage/	IBM Storage home page
www.ibm.com/storage/support	IBM Support home page

Services available and telephone numbers listed are subject to change without notice.

Software Maintenance

All distributed software licenses include Software Maintenance (software subscription and technical support) for a period of 12 months from the date of acquisition providing a streamlined way to acquire IBM software and assure technical support coverage for all licenses. Extending coverage for a total of three years from date of acquisition may be elected. While your Software Maintenance is in effect, IBM will provide you assistance for your 1) routine, short duration installation and usage (how-to) questions; and 2) code-related questions. IBM provides assistance via telephone and, if available, electronic access, only to your information systems (IS) technical support personnel during the normal business hours (published prime shift hours) of your IBM support center. (This assistance is not available to your end users.) IBM provides Severity 1 assistance 24 hours a day, every day of the year.

Getting help online

Be sure to visit the support page for the SAN File System, complete with FAQs, parts information, technical hints and tips, technical publications, and downloadable files, if applicable. This page is at: www.ibm.com/storage/support.

Getting help by telephone

With the original purchase of the SAN File System, you have access to extensive support coverage. During the product warranty period, you may call the IBM Support Center (1 800 426-7378 in the U.S.) for product assistance covered under the terms of the software maintenance contract that comes with SAN File System purchase.

Please have the following information ready when you call:

- SAN File System software identifier, which can be either the product name (SAN File System) or the Product Identification (PID) number
- Description of the problem
- Exact wording of any error messages
- Hardware and software configuration information

If possible, have access to your master console when you call.

In the U.S. and Canada, these services are available 24 hours a day, 7 days a week. In the U.K., these services are available Monday[™] through Friday, from 9:00 a.m. to 6:00 p.m. In all other countries, contact your IBM reseller or IBM marketing representative.¹

1. Response time will vary depending on the number and complexity of incoming calls.

Appendix C. Purchasing additional services

During and after the warranty period, you can purchase additional services, such as support for other IBM and non-IBM hardware, operating systems, and application programs; network setup and configuration; extended hardware repair services; and custom installations. Service availability and name might vary by country.

Appendix D. Disaster recovery

This section points the user to the disaster recovery section of the Admin guide.

Metadata servers

In the event that disaster recovery becomes necessary, you must do a full installation of operating system and the SAN File System on each metadata server engine in the cluster. Then, restore the files that you have previously backed up.

For more detailed information about disaster recovery procedures, refer to the *IBM TotalStorage SAN File System Maintenance and Problem Determination Guide*, found on the publications CD that came with your Metadata servers.

Appendix E. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
MW9A/050
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

The following terms are trademarks of International Business Machines Corporation or Tivoli Systems Inc. in the United States or other countries or both:

AIX	AIX 5L	DB2
Enterprise Storage Server	eServer	FlashCopy
HACMP	IBM	IBM logo

Storage Tank
WebSphere

Tivoli
xSeries

TotalStorage

Intel and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Java and all Java-based trademarks are registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc., in the United States and other countries.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, and service names may be trademarks or service marks

of others.  **JAVA**
COMPATIBLE

Index

Numerics

- 162 configuration error (BIOS) 30
- 184 power password not valid (BIOS) 30

A

- About the Installation and Configuration Guide v
- accessibility
 - disability 123
 - keyboard 123
 - shortcut keys 123
- accessing
 - administrative command-line interface 17
 - SAN File System console 18
- ActiveDirectory 42
- adding a volume to 17
- additional services, purchasing 127
- administrative command-line interface
 - accessing 17
- administrative server package
 - installing 40
- AIX client
 - installation 49
 - installing 50
 - obtaining software 49
 - uninstalling 70
 - validating installation 51
- AIX client, installing 50, 96
- API method
 - backing up 112
- assigning
 - volumes to storage pools 14

B

- backing up
 - API method 112
 - LUN method 111
 - methods 111
- BIOS 30
- boot drive mirror
 - splitting 83

C

- cable
 - routing 26
- cabling
 - external 26
- CD, publications vi, vii
- changes
 - summary of 3
- checklist
 - AIX client upgrade 95
 - installation 19
 - Version 1.1 upgrade 73
 - Windows client 93

- client
 - accessing the global namespace 6
 - creating 118
 - loading the file-system driver 118
 - mounting the global namespace 118
 - setting up 118
- client installation
 - AIX 49
 - Linux 51
 - Solaris 53
 - Windows 47
- client, AIX, installing 50, 96
- clients
 - about fileset 9
- cluster
 - forming 46
 - validating installation 46
- command
 - setupsfs 117
 - setupstclient 118
- commands
 - setautofailover 80
 - stopserver 81
- components
 - SAN File System 5
- configuration error, 162 (BIOS) 30
- configuration files
 - backing up 82
 - restoring 84
- configuring filesets 55
- configuring storage pools 55
- considerations
 - fileset 10
 - nested fileset 11
 - SAN 23
 - security 24
 - zoning 23
- creating
 - client 118
 - fileset for AIX 56
 - fileset for Linux 57
 - fileset for Solaris 57
 - fileset for Windows 58
- creating a policy 64

D

- date and time formats 62
- date and time, setting 34
- DEFAULT_POLICY 59
- Director agent installation 38

E

- Eclipse installation procedure 38

F

- file placement, policy-based 12
- file-placement-rule syntax
 - conventions 59
- fileset
 - about clients 9
 - about metadata servers 9
 - about storage pools 10
 - attaching 8
 - considerations 10
 - considerations for nested 11
 - creating 8
 - creating for AIX 56
 - creating for Linux 57
 - creating for Solaris 57
 - creating for Windows 58
 - creating objects in 8
 - description of 8
 - permissions 11
 - placing in storage pools 8
 - quotas 11
- fileset, configuring 55
- filesets, reassigning 88
- first metadata server
 - preparing for upgrade from version 1.1 80
- FlashCopy Image 115
- functions
 - date and time 62
 - numerical 59
 - string 59

G

- GEE834A 42
- GEJT56A 30
- global fileset
 - description of 8
- global namespace 6
 - client access to 6
 - shared access 8
 - structure of 7

H

- help
 - general 125
 - telephone 125
- high-availability script
 - disabling 80

I

- ibmusbasm installation 39
- installation checklist 19
- installing client for AIX 50, 96

J

Java runtime environment
installation 38

L

label, volume 16
LDAP 24
LDAP, secured 42
Lightweight Directory Access
Protocol 24
limitations 17
limited warranty vi
Linux client
installation 51
installing 52
obtaining software 51
uninstalling 70
validating installation 52
loading the client file-system driver 118
logical unit (LUN) 16
LUN method
backing up 111

M

master metadata server 40
backing up 90
installation 25
preparing for installation 26
software installation 30
upgrading from version 1.1 80
upgrading from version 2.1 103
master role, transferring 87
metadata server
removing from cluster 81
metadata server package
installing 40
Metadata server, uninstalling 69
metadata servers
about fileset 9
metadata volume
limitations 17
migrated data, backing out 67
mounting the global namespace on the
client 118
MPCLI installation 38

N

navigating by keyboard 123
notices 131
notices used in this guide v

O

OpenSLP installation 39

P

package repository
installing 40
package repository, uninstalling 69

permissions
fileset 11
placement policies 59
placement policy rule syntax 59
policies and rules
using 13
policy rules 59
policy set samples 64
policy set, default 59
policy, creating 64
prerequisite software
MPCLI 38
obtaining 30
QLogic driver 35
prerequisite software, clients
SDD 47
prerequisite software, metadata server
Director agent 38
Eclipse 38
ibmusbasm 39
Java runtime environment 38
OpenSLP 39
SDD 39
WebSphere Express 39
publications vi, vii
publications CD vi, vii

Q

QLogic driver
installing 35
obtaining 30
quotas
fileset 11

R

release notes vi, vii
RSA II
configuring 40
disabling watchdogs 81
firmware 42
rule functions 59

S

safety information vi, vii
safety notices, translated vi
sample policy sets 64
SAN considerations 23
SAN File System
components 5
description of 1
major features 2
software 2
SAN File System accessibility
features 123
SAN File System console
accessing 18
SDD installation
clients 47
metadata server 39
security considerations 24
setautofailover command 80
setting time and date 34
setting up 40

setting up the clients 118
setupfs
forming a cluster 46
master metadata server 42
subordinate metadata servers 46
upgrading additional metadata
servers 92
upgrading first metadata server 85
setupfs command 117
setupstclient 118
Solaris client
installation 53
installing 53
obtaining software 53
uninstalling 70
validating installation 54
stopserver command 81
storage management 12
storage pool
assigning volumes to 14
description of 14
system 15
user 16
storage pools 17
about fileset 10
subordinate metadata server
software installation 44
subordinate metadata servers
installing 44
preparing for installation 44
setting up 45
upgrading from version 2.1 103
summary of changes 3
support
general 125
telephone 125
SUSE LINUX Enterprise Server
installing 31
obtaining the software 30
syntax conventions, file-placement-
rule 59
system BIOS 30
system storage pool
description of 15

T

tank.properties 45
time and date, setting 34
time formats 62
trademarks 132
truststore 45

U

uninstalling the Metadata server 69
uninstalling the package repository 69
United Linux service pack
installing 34
obtaining 30
user storage pool
description of 16

V

Version 1.1 upgrade checklist 73

- volume
 - adding to storage pools 17
 - description of 16
 - limitations in the system storage pool 17
- volume label 16
- volumes
 - assigning to storage pools 14
- volumes in the system storage pool 17
- volumes, storage pool 55

W

- watchdogs
 - disabling 81
- Web sites vii, 125
- WebSphere Express installation 39
- Windows client
 - installation 47
 - installing 48
 - obtaining software 47
 - starting 48
 - uninstalling 69
 - validating installation 48

X

- X Window System 34

Z

- zoning considerations 23

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage SAN File System
(based on IBM Storage Tank™ technology)
Installation and Configuration Guide
Version 2 Release 1

Publication No. GA27-4316-01

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>				

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>				
Complete	<input type="checkbox"/>				
Easy to find	<input type="checkbox"/>				
Easy to understand	<input type="checkbox"/>				
Well organized	<input type="checkbox"/>				
Applicable to your tasks	<input type="checkbox"/>				

Please tell us how we can improve this book:

Thank you for your responses. May we contact you? Yes No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

Name

Address

Company or Organization

Phone No.



Fold and Tape

Please do not staple

Fold and Tape



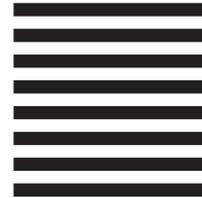
NO POSTAGE
NECESSARY
IF MAILED IN THE
UNITED STATES

BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corp
Dept. CGFA
PO Box 12195
Research Triangle Park, NC 27709-9990



Fold and Tape

Please do not staple

Fold and Tape



Printed in USA

GA27-4316-01

