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

# Installation and Configuration Guide

*Version 2 Release 2*

GA27-4316-02

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

# Installation and Configuration Guide

*Version 2 Release 2*

GA27-4316-02

**Note**
Before using this information and the product it supports, read the information in ″Notices.″

**Third Edition (November 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.

# Contents

# 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 for persons who install and configure SAN File System hardware and software. Those persons 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 is 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 versions 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/sanfs/
- www.ibm.com/storage/software/virtualization/sfs/

The following Web site has information about the languages that have International Components for UNICODE (ICU) converters:
oss.software.ibm.com/cgi-bin/icu/convexp/

## Summary of changes in release 2.2

This section describes the enhancements made to SAN File System in release 2.2.

The following list describes the technical changes and enhancements made to SAN File System for release 2.2.

- **Heterogeneous file sharing** — SAN File System now supports heterogeneous file sharing by implementing user maps that identify equivalent UNIX® and Windows® domain-qualified users.
- **Additional client platforms** — SAN File System supports these additional client platforms:
  - SUSE Linux™ Enterprise Server 8 (32-bit)
- **Installation enhancements** — SAN File System has many installation enhancements, including:
  - **Rolling upgrade** — You can upgrade your metadata servers from SAN File System release 2.1 to 2.2 with minimal disruption in service.
  - **Client upgrade** — Upgrading the SAN File System clients to SAN File System release 2.2 is optional.
- **Management enhancements** — SAN File System has many management enhancements, including:

  - **Non-disruptive file movement** — SAN File System allows you to move a file, along with any FlashCopy® images for that file, to a new storage pool without disruption to the reader.
  - **File defragmentation** — You can defragment an individual file in a storage pool.
  - **File management** — You can create a policy to move files to a new storage pool or delete files to improve the use and balance of premium and inexpensive storage throughout the life cycle of that file.
- **Serviceability enhancements** — SAN File System has many serviceability enhancements, including:
  - **Client tracing** — The tracing utility has been improved to provide a robust, lightweight, buffered, tracing mechanism that is common among all SAN File System client platforms. In addition, log messages are now duplicated in the tracing buffer when tracing is enabled to enhance problem determination.
- **Usability enhancements** — SAN File System has many usability enhancements, including:
  - **Windows client configuration** — You can now use the SAN File System plug-in for the Microsoft® Management Console to configure settings for the Windows clients. This eliminates the need to manually modify the registry.
  - **Metadata checker status indicator** — A progress indicator now shows the progress and estimates the completeness of the metadata checker process.
  - **LUN details** — The LUN details now include the world-wide node name (WWNN) and the world-wide port name (WWPN) associated with each LUN.
  - **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.
  - **Multibyte character set (MBCS)** — SAN File System now supports MBCS in the names of files used in policies and directories in the global namespace to which filesets can be attached:
    - **File names** — Policies accept rules with file names that use MBCS. Rule terms that support MBCS include CHARACTER_LENGTH, CHAR_LENGTH, LEFT, LIKE, POSITION, RIGHT, SUBSTRING and TRIM, and wild cards for zero, one or more characters.
    - **Fileset attach points** — The global namespace now accepts directory names that use MBCS; however, the root fileset attach point (for example, /sanfs) must be in ASCII.

# Chapter 1. Introduction to SAN File System

This topic introduces the 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. 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, SAN File System clients can access data directly from the 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 in which clients 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 to retain its locks.

SAN File System also provides automatic file placement and management through the use of policies and rules. Based on the rules specified in centrally-defined and managed policies, SAN File System automatically stores, moves, and deletes data 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 platform 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.

**Heterogeneous file sharing**

All clients, regardless of operating system or hardware platform, have uniform access to the data stored (under the global namespace) in SAN File System. File metadata (such as last modification time) is 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 and management of files on appropriate storage devices. You can define storage pools based on specific application requirements and quality of services, and define rules based on data attributes to store the files on the appropriate storage devices automatically. SAN File System provides 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 these 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* is a set of CDs that contain the master console software and documentation.

## Components

Figure 1 on page 3 illustrates the major components of SAN File System.

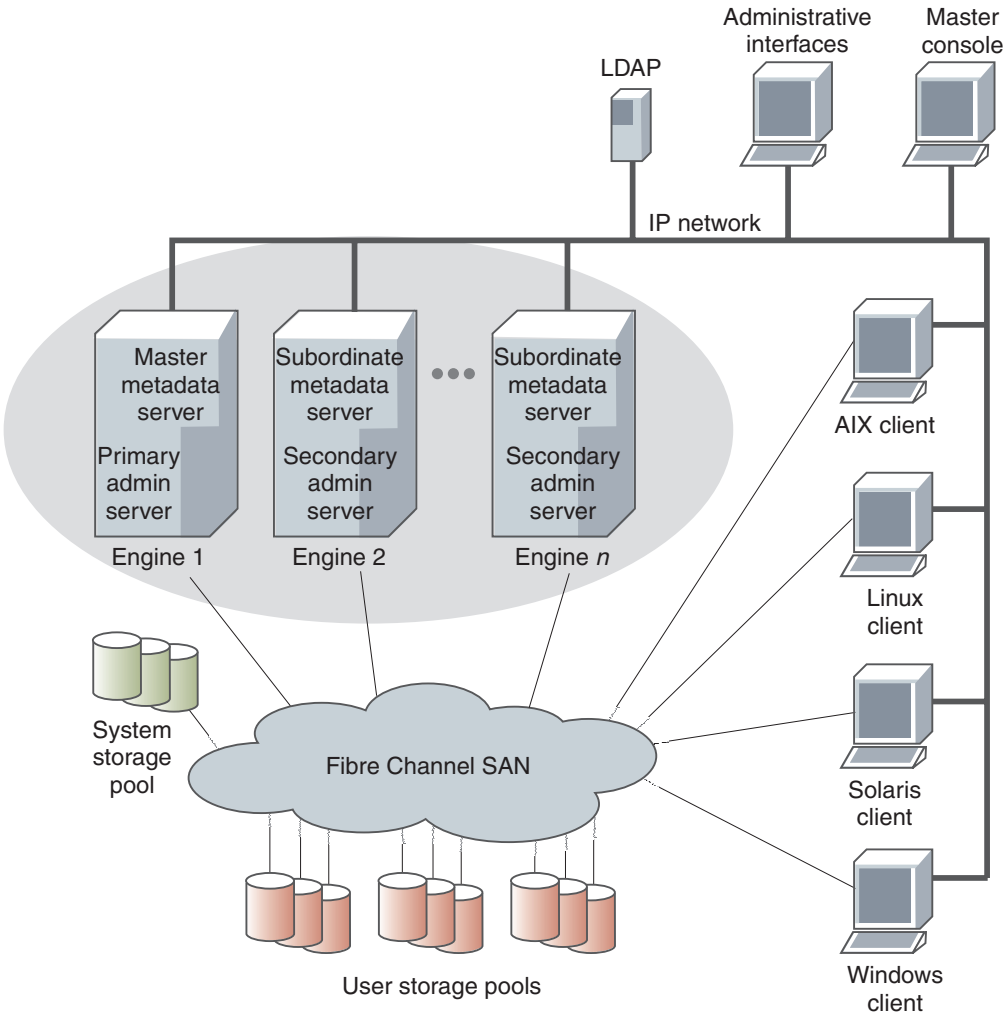


*Figure 1. SAN File System components*

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 machine that you use to access these interfaces is called the *administrative console*. The administrative server uses a *Lightweight Directory Access Protocol (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 share data and have their storage centrally managed by SAN File System 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 for 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

## Engines

Within SAN File System, the hardware on which the metadata servers and administrative servers run are called storage *engines*. SAN File System supports from two to eight engines.

SAN File System is intended to run with a minimum of two engines; however, you can run a single-engine system if:

- All of the other engines fail (for example, if you have only two engines, and one of them fails)
- You want to bring down all of the engines except one before performing scheduled maintenance.
- One engine hosts a spare metadata server.

You can use the SAN File System console or administrative command-line interface to monitor and control the engines from any computer with a network connection to the cluster.

## 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 and 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). Filesets divide the global namespace into a logical, organized structure. They 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 is aware of 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 the 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 other than the root fileset, it is called a *nested fileset*.

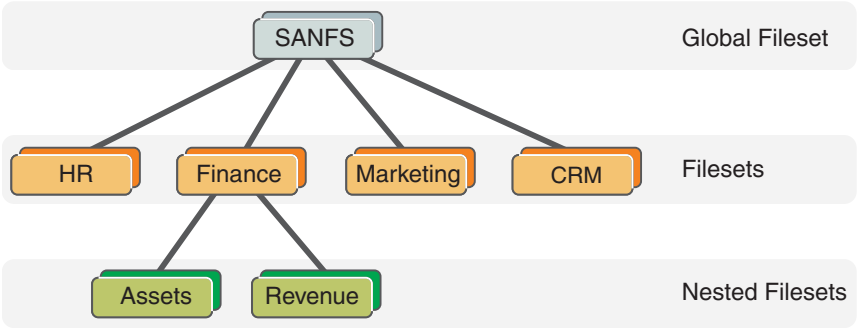


You can detach a fileset and reattach it at the same location or a different location. If a fileset is reattached at a different location, all the files contained in the fileset are rooted to the new location without any further operations. Before a fileset can be detached, any nested filesets must be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

A client cannot create hard links across fileset boundaries and cannot rename, move, or delete a directory that is the root of a fileset. If a client attempts to perform any of these operations, SAN File System returns an error indicating a cross-file-system condition.

## Filesets and metadata servers

When creating a fileset, you can statically assign the fileset to a specific metadata server or SAN File System can dynamically assign it to a metadata server for you. Filesets that are statically assigned are known as *static filesets*. Filesets that are dynamically assigned are known as *dynamic filesets*.

The assigned 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 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. Figure 2 on page 7 shows an example of the relationship between filesets and storage pools.

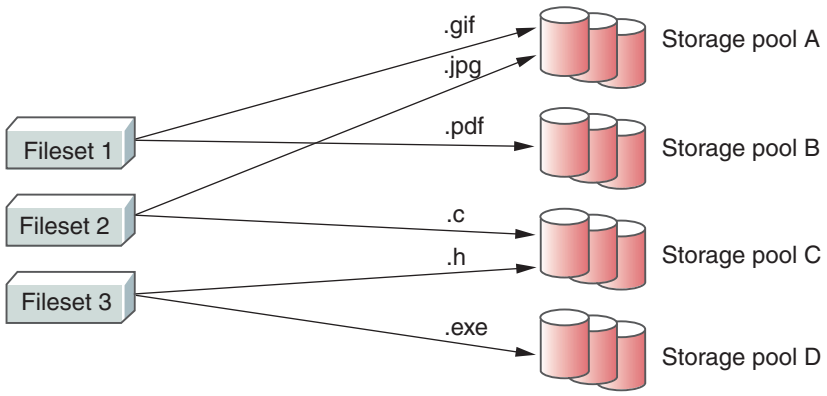


*Figure 2. The relationship between filesets and storage pools*

Using file-placement policies, you can specify that all files created in a particular fileset are to be stored in a specific storage pool. Using file-management policies, you can define how files in a specific fileset are to be moved or deleted during the file's life cycle.

## 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 criteria when placing individual files in global namespace.

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. Also, when the number of filesets is greater than one thousand, response time will increase when you issue fileset commands.

To facilitate file sharing, you can optionally separate filesets by their *primary allegiance* of the operating system. 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.

## 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 the 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 limit*, 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%, an alert is generated when the fileset contains 80 GB of data. (Note that the quota is based on space allocated to the fileset, not the data is contains.)

The action taken when the fileset reaches its quota size depends on whether the quota is defined as hard or soft. If you use a hard quota, once the threshold is reached, SAN File System denies new client requests to add more space to the fileset (by creating or extending files). If you use a soft quota, which is the default, SAN File System allocates more space but continues to send alerts. Once the amount of physical storage available to global fileset is exceeded, no more space can be used. You can set the quota limit, threshold and quota type 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 nested filesets.
- 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 parent fileset becomes a rogue fileset and is unable to be failed over, then the nested filesets of that parent fileset would also, effectively, be unavailable.
- 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 FlashCopy 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 is more complex.

## Global namespace

The *global namespace* is the key to SAN File System. It gives all SAN File System clients common access to all files and directories, 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 the 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.

Figure 3 on page 10 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 file system, 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.

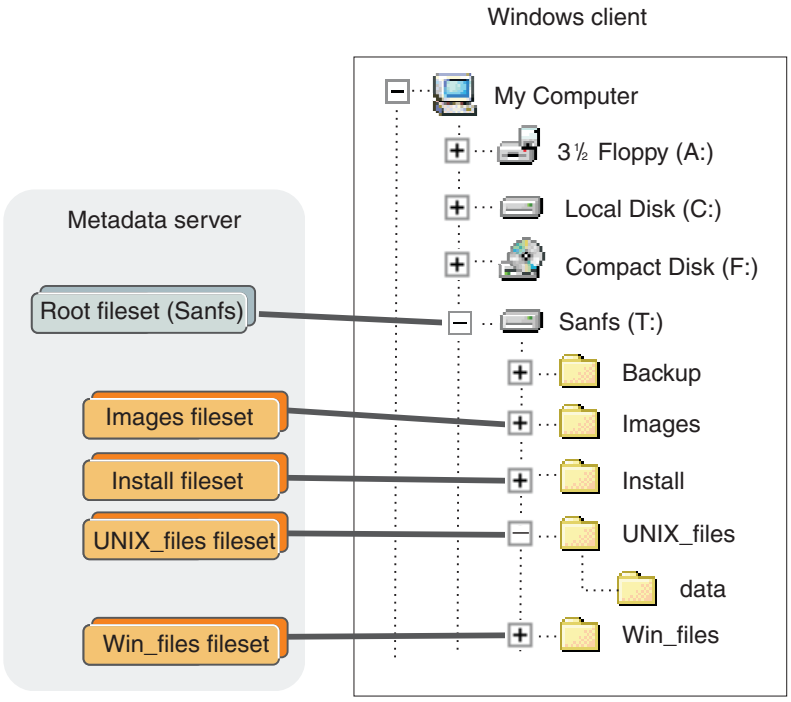


*Figure 3. Filesets as seen by the metadata server and client*

## 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).

Figure 4 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.

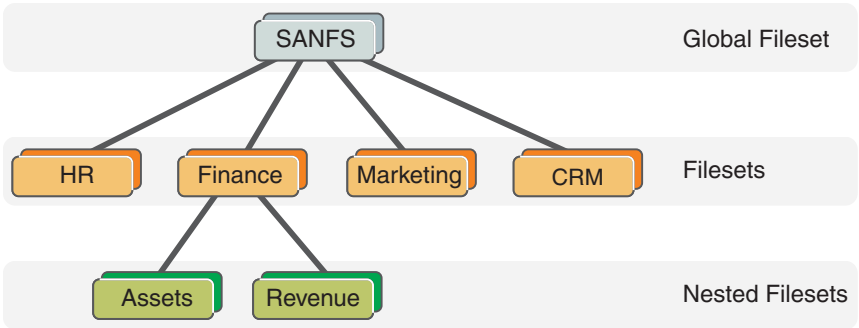


*Figure 4. Sample global namespace*

### 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.

## Storage management

SAN File System provides automatic file placement and management through the use of policies. The policy *rules* cause newly created files to be placed in the appropriate storage pools and cause files of a certain age or size to be moved or deleted.

### 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 place newly created files into the appropriate storage pool automatically.

The file-placement policy tells a metadata server where to place the data for a newly created file in a specific storage pool if the attributes of that file meet the criteria specified in a rule. 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. Other criteria include these:

- Date and time when the file is created
- Fileset
- File name or extension
- User ID and group ID on UNIX clients

The rules in a file-placement policy are evaluated in order until the condition in one of the rules is met. The data for the file is then stored in the storage pool that is specified by the applicable rule. If none of the conditions specified in the rules is met, the data for the file is stored in the default storage pool.

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 (CLI). You can also use the **statpolicy** command from the administrative CLI to view the statistics about the file-placement policy rules.

**Attention:**

It is recommended that you do not use creation time, user ID or group ID to place file. If you do base any file-placement rules on creation time, user IDs, or group IDs, be aware of these restore and migration considerations:

- A rule that uses the creation date as the placement criteria assigns a file based on the time that the file was restored or migrated, not the original creation time.
- A rule that uses a user ID or group ID as the placement criteria assigns a file based on the effective user and group IDs of the restore process, not the original file's user and group IDs.

## Policies and rules

This topic describes how SAN File System automates the management of files using policies and rules.

SAN File System enables you to automate the management of files using policies and rules. Properly managing your files allows you efficiently use and balance your premium and inexpensive storage. SAN File System supports these policies:

- File-placement policies are used to automatically place newly created files to a specific storage pool.
- File-management policies are used to manage files (move or delete) during its lifecycle by moving them to another storage pool or delete them all together.

**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. Manually 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 file-placement 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.

Currently, there is no error checking for file-management policies.

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 file-placement policy is created and remains active until you create and activate a new one. The default file-placement 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 what to do with 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 identify the conditions, such as these, that when matched causes that rule to be applied:
- Date and time when the file is created
- Date and time when the file was last accessed
- Fileset
- File name or extension
- File size
- 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 what is to be done with that file. For example, when a client creates a file, SAN File System scans the list of rules in the active file-placement 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.

**Attention:**

It is recommended that you do not use creation time, user ID or group ID to place file. If you do base any file-placement rules on creation time, user IDs, or group IDs, be aware of these restore and migration considerations:
- A rule that uses the creation date as the placement criteria assigns a file based on the time that the file was restored or migrated, not the original creation time.
- A rule that uses a user ID or group ID as the placement criteria assigns a file based on the effective user and group IDs of the restore process, not the original file's user and group IDs.

## Storage pools

 Watch and learn

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 assign 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 Figure 5) 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.

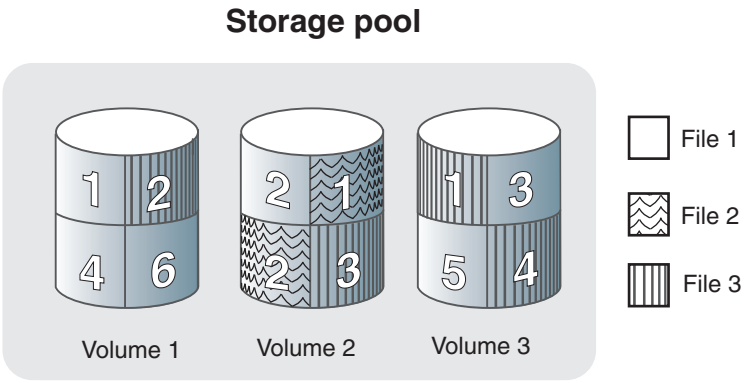


*Figure 5. File space allocation*

**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 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.

**Restriction:** A metadata server can access up to a combined total of 256 SCSI disk single-pathed and/or vpath multi-pathed LUNs. This is a limitation of the Linux operating system. When the number of entries in the storage subsystem device driver's vpath configuration file reaches 256, 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 the 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 user interfaces for managing the SAN File System.

## Accessing the administrative CLI

This topic describes how to start the administrative command-line interface (CLI).

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

1. Log in directly to the engine, or from another workstation through Secure Shell (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. If you have not already set the password, you can do this in one of two ways:
   - Set the password using the tankpasswd utility.
   - Set the SFS_CLI_PASSWDFILE environment variable to the location of the password file.
3. Enter the **sfscli** command to start the administrative CLI to run commands in interactive mode.

## Accessing the SAN File System console

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

If you are running Microsoft Windows 2003 operating system, you must first add the metadata server to your list of Trusted Sites before you can access the SAN File System console.

After you have added the metadata server to the list of Trusted Sites (or when you are running an operating system other than Windows 2003), you can access the SAN File System console by following these steps:

1. In a browser window, enter `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 is displayed, enter your administrator ID and password.

# Chapter 2. Pre-installation tasks

This topic provides a high-level overview of the procedures for installing SAN File System in your environment. It also includes the steps you take to prepare your environment for installation.

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. You need a licensed copy of SUSE Linux Enterprise Server 8 for each of the metadata server engines in the cluster. For more information about obtaining SUSE Linux, 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" |
| | a | Prepare the SAN. | "SAN considerations" on page 20 |
| | b | Prepare switch zoning. | "Zoning considerations" on page 20 |
| | c | Prepare storage devices. | "Supported storage subsystems" on page 21 |
| | d | | |
| 2 | Prepare LDAP. | | "LDAP configuration" on page 22 |
| 3 | Prepare the engine for installation. | | "Preparing the engine for installation" on page 41 |
| | a | Cable the metadata server. | "Cabling" on page 42 |
| | b | Obtain software prerequisites. | "Obtain prerequisite software" on page 46 |
| | c | Upgrade system BIOS, if needed. | "Upgrading system BIOS" on page 46 |
| | d | Mirror boot drives. (This step applies only when you are using an LSI SCSI RAID controller.) | LSI Logic Configuration Program documentation |

## 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:

- SAN considerations
- Zoning considerations
- Security considerations (LDAP configuration)
- Supported storage subsystems

For additional information about these considerations, refer to the *SAN File System Planning Guide*, which also contains planning worksheets.

## 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 are 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:

- Because of 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 protects 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).

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

## 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 need to add SAN File System administrative 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 TotalStorage DS4300 Turbo (IBM TotalStorage FAStT600T), DS4400 (FAStT700), and DS4500 (FAStT900) running firmware version 8.4 on the storage device and software version 8.41 on the client platforms

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.

**Note:** If you are using a DS4000 Series storage subsystem, all DS4000 LUNs that are used in the system storage pool must be in a separate partition from the LUNs that are used by the user storage pools.

Refer to the IBM TotalStorage 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 are used in your SAN File System environment to ensure that they are all compatible.

DS4000 LUNs assigned to user storage pools can be shared by multiple SAN File System clients as long as the clients are aware of one operating system type. LUNs within a DS4000 partition can only be used by one operating system. If you use DS4000 firmware lower than version 8.41, you cannot have more than 32 LUNs per partition.

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.

## LDAP configuration

The SAN File System administrative agent relies on your LDAP installation to authenticate and authorize each administrative operation based on your authentication model. This support requires that the LDAP service be readily available when an administrator command is issued. If the LDAP service cannot be reached, the administrative operation fails with an authentication error. Therefore, you need to ensure that your LDAP service has high availability.

SAN File System requires some configuration of the LDAP server to use LDAP to authenticate SAN File System administrators. In general, this configuration requires that you provide the following types of information:

**Network**
You must identify the machine on which the LDAP server is running (and port if not running in the default port normally used by the LDAP server). SAN File System authorization of the LDAP server also requires an authorized LDAP user name. The authorized LDAP user name must be able to browse the LDAP tree where the users and roles are stored.

**Users** All administrative users must have an entry in the LDAP database. They must have the same objectClass and make use of one attribute to consistently store the administrator's login user ID. They must contain a "user ID" type of attribute.

**Roles** Each of the roles that you plan to use must have an entry in the LDAP database. These roles must have the same parent distinguished name (DN). Each must have an attribute containing the string that describes its role:

Administrator, Backup, Operator, or Monitor. Finally, each must support an attribute that can contain multiple values, one value for each DN of the role occupant.

You define these four roles in the LDIF file. You can change the default values of these roles to values that are unique to your organization.

You can use the worksheet in Table 1 to compile this information. You also need to import an LDIF file.

The procedures to set up three possible LDAP infrastructures are provided as examples of how to configure LDAP. Those examples might differ from your configuration, but can provide some helpful guidance. The three LDAP configuration alternatives are:

- "Configuring LDAP using IBM Directory Server Version 5.1" on page 25
- "Configuring LDAP using OpenLDAP" on page 29
- "Configuring LDAP using Microsoft Active Directory LDAP" on page 35

*Table 1. LDAP planning worksheet*

| Description | Example value | Your value |
|---|---|---|
| IP address | 9.42.164.125 | |
| Port numbers | 389 insecure; 636 secure | |
| Authorized LDAP user name | cn=root | |
| Authorized LDAP password | *secret* (default for IBM Directory Server) | |
| Attribute containing login user ID | uid for IBM Directory Server and OpenLDAP; sAMAccountName for MS Active Directory | |
| Role parent DN | dn: ou=Roles, o=yourOrg objectclass: organizationalUnit | |
| Attribute containing role name | cn | |
| Attribute for role occupants | roleOccupant for IBM Directory Server and OpenLDAP; description for Microsoft Active Directory | |

## LDIF file

LDAP configurations are specified in a format known as LDAP Data Interchange Format (LDIF).

**Purpose**

This text-based file stores information in object-oriented hierarchies of entries. LDIF is used to import and export directory information between LDAP-based directory servers, or to describe a set of changes that are to be applied to a directory. The purpose of using an LDIF file is that you can populate the LDAP directory using a single file rather than having to populate the directory one entry at a time.

The process for configuring the IBM Directory Server calls for importing a LDIF file, so if you are installing the IBM Directory Server, you need to create and save

the LDIF file before you begin the IBM Directory installation process.

**Sample**

You can use the sample LDIF configuration file shown here as-is, with the default values, or you can customize the file with values that are unique to your organization. To do this, copy-and-paste the contents of the sample file into a text file and with a text editor complete the following steps:

1. Replace all occurrences of the string *yourOrg* with an appropriate value for your company or department.
2. Replace the Administrator role user ID (*Admin*) and password (*adminpassword*) with appropriate values for your company or department.
3. Replace the Backup role user ID (*Back*) and password (*backpassword*) with appropriate values for your company or department.
4. Replace the Operator role user ID (*Oper*) and password (*operpassword*) with appropriate values for your company or department.
5. Replace the Monitor role user ID (*Monit*) and password (*monitpassword*) with appropriate values for your company or department.
6. If you require additional users (such as a second Administrator), you can copy and paste the rows that define the particular user, and make the appropriate edits. For example, if you need an additional Administrator, you can edit the values for the attributes uid (Admin) and userPassword (adminpassword), as highlighted in this example (the other attributes would be customized for your company or department too):

   ```
   dn:cn=yourOrgAdmin Administrator,ou=Users,o=yourOrg
   objectClass:inetOrgPerson
   cn:yourOrgAdmin Administrator
   sn:Administrator
   uid:Admin
   userPassword:adminpassword
   ```

7. Save the customized file as a text file to a directory that can be accessed from the server on which the LDAP directory is installed.

```
dn:o=yourOrg
objectClass:organization
o:yourOrg

dn:ou=Users,o=yourOrg
objectClass:organizationalUnit
ou:Users

dn:cn=yourOrgAdmin Administrator,ou=Users,o=yourOrg
objectClass:inetOrgPerson
cn:yourOrgAdmin Administrator
sn:Administrator
uid:Admin
userPassword:adminpassword

dn:cn=yourOrgMon Monitor,ou=Users,o=yourOrg
objectClass:inetOrgPerson
cn:yourOrgMon Monitor
sn:Monitor
uid:Monit
userPassword:monitPassword

dn:cn=yourOrgBack Backup,ou=Users,o=yourOrg
objectClass:inetOrgPerson
cn:yourOrgBack Backup
sn:Backup
uid:Back
userPassword:backpassword
```

```
dn:cn=yourOrgOper Operator,ou=Users,o=yourOrg
objectClass:inetOrgPerson
cn:yourOrgOper Operator
sn:Operator
uid:Oper
userPassword:operpassword

dn:ou=Roles,o=yourOrg
objectClass:organizationalUnit
ou:Roles

dn:cn=Administrator,ou=Roles,o=yourOrg
objectClass:organizationalRole
cn:Administrator
roleOccupant:cn=yourOrgAdmin Administrator,ou=Users,o=yourOrg

dn:cn=Monitor,ou=Roles,o=yourOrg
objectClass:organizationalRole
cn:Monitor
roleOccupant:cn=yourOrgMon Monitor,ou=Users,o=yourOrg

dn:cn=Backup,ou=Roles,o=yourOrg
objectClass:organizationalRole
cn:Backup
roleOccupant:cn=yourOrgBack Backup,ou=Users,o=yourOrg

dn:cn=Operator,ou=Roles,o=yourOrg
objectClass:organizationalRole
cn:Operator
roleOccupant:cn=yourOrgOper Operator,ou=Users,o=yourOrg
```

## Configuring LDAP using IBM Directory Server Version 5.1

This topic lists the overall steps that must be performed when configuring an LDAP server using IBM Directory Server Version 5.1 with SAN File System.

IBM Directory Server is a powerful LDAP infrastructure that is designed to provide a foundation for deploying comprehensive identity management applications and advanced software architectures. To use IBM Directory Server with SAN File System, you need to complete the following tasks:

- Install IBM Directory Server
- Create the LDAP database
- Configure IBM Directory Server for SAN File System
- Start the LDAP server and configure the Administrative server

### Installing IBM Directory Server Version 5.1

This procedure lists steps that you can use as examples to configure your LDAP server when using IBM Directory Server Version 5.1.

1. Download the IBM Directory Server Version 5.1 software from the Web site: www14.software.ibm.com/webapp/download/search.jsp?rs=ldap&go=y.
   a. Choose an appropriate platform.
   b. Select a language.
   c. Register.
   d. Define a user ID and password, and then fill in the questionnaire and accept the license agreement.
   e. Choose the CD-image or zip-file download. The CD-image download is copied to a temporary directory. From that temporary directory, copy the image onto a CD. When you insert the CD into the drive, the installation wizard starts automatically.

If you download the zip file, unzip it and run setup.exe.

f. After you download either the CD image or the zip file and start the installation, continue with step 2.

2. Select a language to use for the installation and click **OK**.
3. Click **Next** to continue when the Welcome screen appears.
4. Select the button that indicates that you accept the terms when the license agreement appears. Then click **Next**.
5. Select a directory to install IBM Directory Server. The default is C:\Program Files\IBM\LDAP. You can accept this or enter an alternative, and then click **Next**.
6. Select the language for IBM Directory Server and click **Next**.
7. Select the setup type. This example uses Typical. Click **Next**.
8. Accept the defaults in the features window.
9. Specify a user ID and password for IBM DB2®. DB2 is used as the underlying repository and is installed automatically. You can specify a new user ID or an existing one. If the ID exists, then the password must be correct.
10. Click **Next** if all of the installation options are correct in the summary of the installation options that appears.
11. Click **OK** to continue when a window appears that states that DB2 will install in the background. This installation can take up to 20 minutes.
12. After some time, an IBM Global Security Toolkit (GSKit) message appears. The GSKit installs in the background, which can take up to 5 minutes. The GSKit provides a Secure Sockets Layer (SSL) with encryption strengths up to triple data encryption service (DES). A DOS window is displayed.
13. After some time, the IBM WebSphere® Application Server Express window appears. This application installs in the background, which can take up to 10 minutes. WebSphere provides the application environment for IBM Directory Server.
14. After this installation is complete, the readme file for the IBM Directory Server client is displayed. Review it and click **Next** to continue.
15. The server readme file is displayed. Review it and click Next to continue.
16. You are prompted to restart your system now or later. A reboot is required to complete the installation. Select *Yes, restart my system* and click **Next**. The installation complete window opens. Click **Finish** to continue and reboot the server.

Continue with *Creating the LDAP database*.

## Creating the LDAP database

This topic provides the steps necessary to create the LDAP database when you are using IBM Directory Server Version 5.1 as your LDAP infrastructure.

You must have completed *Installing IBM Directory Server Version 5.1* before completing this procedure.

After the system has rebooted after installing the IBM Directory Server software, the IBM Directory Server configuration tool starts automatically. If it does not start, launch it using **Start → Programs → IBM Directory Server 5.1 → Directory Configuration**.

1. Click **Administrator DN/password** in the left column to open the configuration tool in the introduction panel.

2. Set the IBM Directory Server's Administrator DN and password. This example uses the following values:
   - Administrator DN: *cn=Manager,o=yourOrg*
   - Password: *password*

   Click **OK** to continue.

   **Note:** The IBM Directory Server's Administrator is the administrator of the LDAP directory, and is not the actual entry in the directory. The IBM Directory Server's Administrator should not be confused with an administrator listed in the example LDIF file.
3. Click **OK** when a confirmation appears indicating that the Administrator DN and password have been successfully set.
4. Click **Configure database** in the left column.
5. Select **Create a new database** in the Configure Database panel and click **Next**.
6. When prompted, enter the user ID for the DB2 database that was specified during the installation and click **Next**.
7. Enter a name for your LDAP database and click **Next**. This database is created in DB2.
8. Specify the code page for your DB2 database by selecting the default **Create a universal DB2 database**, and click **Next**.
9. Select the drive letter on which to create the LDAP database and click **Next**. The default is the C partition.
10. Verify the entries on the summary screen and click **Finish** to create the database.
11. When the database is created, click **Close**. A series of task messages is displayed in the IBM Directory Server Configuration Tool panel to inform you of configuration status.

Continue with *Configuring IBM Directory Server for SAN File System*.

## Configuring Directory Server for SAN File System

This topic provides the steps required to configure IBM Directory Server Version 5.1 in order to use it with SAN File System.

You must have completed *Installing IBM Directory Server Version 5.1* and *Creating the LDAP database* before starting this procedure.

After you install IBM Directory Server 5.1 and create the database, you need to create and import your configuration.

LDAP configurations are specified in a format known as LDAP Data Interchange Format (LDIF). You can see a sample LDIF file in "LDIF file" on page 23.

1. Start the IBM Directory Server configuration tool: **Start → Programs → IBM Directory Server 5.1 → Directory Configuration**.
2. Click **Manage Suffixes**.
   a. In the Suffix DN field, type *o=yourOrg*.

      Note that *o=yourOrg* is the default value provided in the example LDIF. If you edited this value, make sure your entry matches the edited value.
   b. Click **Add**.
   c. Click **OK**.
3. Click **Import ldif data** in the left column.

4. Enter the file name of your saved LDIF configuration file.

   **Tip:** IBM Directory Server expects a c:\tmp directory on your system drive when importing an LDIF file. Make sure that you have this directory; if it does not exist, create it.
5. Click **Import** to start importing the LDIF file.
6. Close the configuration tool when the import completes. Your LDAP server is now configured to be used with SAN File System.

Continue with *Configuring the administrative server*.

## Configuring the administrative LDAP server

This topic lists the steps required to set up access to the administrative server that can be found on the LDAP server.

You must complete *Installing IBM Directory Server Version 5.1*, *Creating the LDAP database*, and *Configuring Directory Server for SAN File System* before completing this procedure.

1. Start the LDAP server.
2. Change directories to *C:\Program Files\IBM\LDAP\bin*
3. Enter the **ibmslapd** command to start the Directory Server. The Directory Server starts, and you see the following status messages:

   **Note:** When the Directory Server starts, leave this window open in the background. If you close the window, the Directory Server stops.

```
C:\Program Files\IBM\LDAP\bin>ibmslapd
Server starting.
Plugin of type EXTENDEDOP is successfully loaded from
        libevent.dll.
Plugin of type EXTENDEDOP is successfully loaded from
        libtranext.dll.
Plugin of type EXTENDEDOP is successfully loaded from
 libldaprepl.dll.
Plugin of type PREOPERATION is successfully loaded from
 libDSP.dll.
Plugin of type EXTENDEDOP is successfully loaded from
 libevent.dll.
Plugin of type EXTENDEDOP is successfully loaded from
 libtranext.dll.
Plugin of type AUDIT is successfully loaded from C:/Program
 Files/IBM/LDAP/bin/l
 ibldapaudit.dll.
Plugin of type EXTENDEDOP is successfully loaded from
 libevent.dll.
Plugin of type EXTENDEDOP is successfully loaded from
 libtranext.dll.
Plugin of type DATABASE is successfully loaded from C:/Program
 Files/IBM/LDAP/bi
n/libback-rdbm.dll.
Plugin of type REPLICATION is successfully loaded from C:/Program
 Files/IBM/LDAP
/bin/libldaprepl.dll.
Plugin of type EXTENDEDOP is successfully loaded from
 libevent.dll.
Plugin of type DATABASE is successfully loaded from C:/Program
 Files/IBM/LDAP/bi
n/libback-config.dll.
Error code -1 from odbc string:" SQLConnect " SFSLDAP .
SQL1063N DB2START processing was successful.
Plugin of type EXTENDEDOP is successfully loaded from
```

```
 libloga.dll.
Non-SSL port initialized to 389.
IBM Directory (SSL), Version 5.1 Server started.
```

4. From another command prompt, change directories to *C:\Program Files\IBM\ldap\appsrv\bin*.
5. Enter the **startserver.bat server1** command to start the administrative server. The administrative server starts, and you see the following status messages:

```
C:\Program Files\IBM\LDAP\appsrv\bin>startserver.bat server1
ADMU0116I: Tool information is being logged in file C:\Program
Files\IBM\LDAP\appsrv\logs\server1\startServer.log
ADMU3100I: Reading configuration for server: server1
ADMU3200I: Server launched. Waiting for initialization status.
ADMU3000I: Server server1 open for e-business; process id is 1916
```

6. Close the command prompt after the administrative server starts.
7. Verify that the administrative server is operating correctly:
   a. With a browser, view the IBM Directory Server at (replacing *localhost* with your host name)

      `http://localhost: 9080/IDSWebApp/IDSjsp/Login.jsp`

   b. Enter the default user name *superadmin* and password *secret* and click **Login**. The main administrator console is displayed.
8. Change the default administrator login password, if necessary:
   a. Select **Change console administrator login** from the left column.
   b. Enter a new password and click **OK**.
9. Click **Manage console servers**.
10. Add the host name of your local machine.
11. Click **Add** and the Add server panel is displayed.
12. Enter the host name (the short name or full name) and leave the other options as their default. Click **OK**.
13. The list of console servers is displayed, showing that the local host is added.
14. You can ensure that the local host has been added correctly by logging out and then logging in again, using your host name and LDAP user name.
    a. Select the localhostname in the LDAP Hostname drop-down list.
    b. Enter the user name and password as defined in the LDIF file that you imported.
    c. Click **Login**.
15. The default administrator console, IBM Directory Server Web Administration Tool, is displayed.

IBM Directory Server is now installed and ready for use by SAN File System.

## Configuring LDAP using OpenLDAP

This topic lists the overall steps that are required when configuring an LDAP server using OpenLDAP.

OpenLDAP software is an open-source implementation of LDAP. This topic lists the steps to install OpenLDAP 2.0.x on Red Hat Linux. Specific versions of LDAP are required, depending on the version of Red Hat Linux. Table 2 on page 30 shows the Red Hat version and the OpenLDAP build versions.

*Table 2. OpenLDAP build versions*

| Red Hat Linux release | OpenLDAP build version |
|---|---|
| AS2.1 | 2.0.21 |
| 7.3 | 2.0.23 |
| 8.0 | 2.0.25 |
| 9.0 | 2.0.275 |

The overall steps required are:

1. Install the OpenLDAP packages
2. Configure the OpenLDAP client
3. Configure the OpenLDAP server
4. Configure OpenLDAP for SAN File System

The instructions are based on Red Hat Linux 9.0, but other versions should be similar. The instructions for SUSE Linux Enterprise Server 8 should also be similar.

## Install OpenLDAP packages

This topic provides the steps required to install the OpenLDAP packages when you plan to use OpenLDAP as your LDAP architecture.

1. Determine the Red Hat Linux release that is installed. The release number is stored in the /etc/redhat-release file. If you are running SUSE Linux, then the release number is in the /etc/SUSE-release file.
2. Determine the version of OpenLDAP that is currently installed by entering the **rpm -qa | grep openldap** command at the Linux prompt:

   ```
   # rpm -qa |grep openldap
   openldap-2.0.27-8
   openldap-clients-2.0.27-8
   openldap-servers-2.0.27-8
   ```

   If a default Red Hat Linux installation was used, there is at least one OpenLDAP Red Hat Package Manager (RPM) installed. If you have the correct version installed, continue with "Configure the OpenLDAP client" on page 31. If no OpenLDAP RPMs are installed or an RPM is not a valid version, you need to install the correct version.

   The required RPMs for an LDAP server on Red Hat Linux are openldap-2.0.xx-F, openldap-client-2.0.xx-F, and openldap-server-2.0.xx.F, where 2.0.xx-F corresponds to Table 3.

*Table 3. OpenLDAP versions*

| Red Hat Linux release | OpenLDAP build version |
|---|---|
| AS2.1 | 2.0.21 |
| 7.3 | 2.0.23 |
| 8.0 | 2.0.25 |
| 9.0 | 2.0.275 |

3. The LDAP RPMs can either be found on the Red Hat CD or downloaded from one of the following RPM download sources:
   - www.rpmfind.net, search on openldap and select the RPM based on the distribution.

- www.redhat.com, select **Download**, and then search on openldap. Other distributions might not be listed here.

**Tip:** You only need to download RPMs that are not installed. For example, if you have openldap-2.0.xx and openldap-client-2.0.xx installed but not openldap-server-2.0.xx, then you only need to download the openldap-server-2.0.xx package.

4. After downloading the RPMs to the Linux server, change to the download directory and start the installation using the **rpm** command.

```
# rpm -ivh openldap*
```

The RPMs are installed, with a hash-mark progress bar. If the RPMs are not installed due to any missing prerequisite RPMs, find the RPMs using step 3 on page 30. If, however, the RPMs do not install because of missing prerequisite files or mismatched file versions, the RPM version selected is not appropriate for the Red Hat Linux installation. Investigate the specific files in conflict, and confirm which OpenLDAP RPM version matches those files.

5. Verify that the OpenLDAP RPMs have been installed with the **rpm -qa | grep openldap** command at the Linux prompt.

```
# rpm -qa | grep openldap
openldap-2.0.27-8
openldap-clients-2.0.27-8
openldap-servers-2.0.27-8
```

The three RPMs are installed: the base, the clients, and the servers.

Continue with *Configure the OpenLDAP client.*

## Configure the OpenLDAP client

This topic outlines how to configure the OpenLDAP client. This component is not required for SAN File System but is included for reference if you need to extend your LDAP configuration.

Complete *Install OpenLDAP packages* before starting these steps.
Customize the client configuration file /etc/openldap/ldap.conf. This file contains information for the LDAP clients. It holds default values so that all values do not have to be specified on the command line:

```
# $OpenLDAP: pkg/ldap/libraries/libldap/ldap.conf,v 1.4.8.6
2000/09/05 17:54:38 kurt Exp $
#
# LDAP Defaults
#
# See ldap.conf(5) for details
# This file should be world readable but not world writable.
#BASE dc=example, dc=com
#URI ldap://ldap.example.com ldap://ldap-master.example.
com:666
#SIZELIMIT 12
#TIMELIMIT 15
#DEREF never
HOST 127.0.0.1
BASE o=yourOrg
```

1. Edit the ldap.conf file using a text editor such as vi. Table 4 on page 32 defines the values to be customized for your installation.

*Table 4. ldap.conf file parameters*

| Parameter | Description |
|---|---|
| Host | Set this parameter to the IP address of the host used by the local LDAP clients. |
| Base | This parameter is the base DN for any searches. Searches such as those performed by the **ldapsearch** command are restricted to this DN by default. In our example, it is o=yourOrg. This parameter should reflect the root of your LDAP tree. |

2. Save and quit from the editor. For a more detailed description of this file, refer to the manual page (using the **man ldap.conf** command).

Continue with *Configure the OpenLDAP server*.

## Configure the OpenLDAP server

This section shows how to configure the OpenLDAP server.

Complete *Install OpenLDAP packages* and *Configure the OpenLDAP client* before starting this procedure.

The server is known as the stand-alone LDAP daemon (or slapd).

1. Edit the entries in the configuration file /etc/openldap/slapd.conf. These entries are shown in Table 5.

*Table 5. slapd.conf parameters*

| Parameter | Description |
|---|---|
| suffix | This parameter is the base DN for any searches. Searches such as those set this parameter to the base suffix specified in Table 3. In our example, it is o=yourOrg. |
| rootdn | This parameter is the DN of the LDAP "root" user. Although it can have any hierarchy, it can most easily be placed under the suffix. In our example, it is cn=Manager,o=yourOrg. |
| rootpw | This parameter is set to a shielded (not encrypted) password at the next step. |

Here is an example slapd.conf file:

```
#suffix "dc=my-domain,dc=com"
suffix "o=yourOrg"
#rootdn "cn=Manager,dc=my-domain,dc=com"
rootdn "cn=Manager,o=yourOrg"
#
# Cleartext passwords, especially for the rootdn, should
# be avoided. See slappasswd(8) and slapd.conf(5)
for details.
# Use of strong authentication encouraged.
# rootpw secret
```

2. Save the slapd.conf file and quit.
3. Create a shielded password for the root DN. Enter the following command with the parameter **slappasswd** enclosed in back-slanting tick marks (`slappasswd `):

```
# export SLAPPW=`slappasswd`
```

4. When prompted, enter the same password twice. It is concealed as is other UNIX password input.

   **Note:** The slappasswd utility creates an encrypted Switch Link Authentication Protocol (SLAP) password. If you enter the command **echo $SLAPPW**, you see the encrypted password string, starting with the characters *{SSHA}*, followed by apparently random characters. The {SSHA} prefix indicates that SSHA has been used as the default encryption algorithm.
5. When you return to the prompt, the *SLAPPW* variable contains the shielded string that is needed for the slapd.conf file. Insert the value of this variable into the slapd.conf file. Be careful to enter this string exactly, especially if you are not familiar with Linux command syntax:

   ```
   # echo "rootpw $SLAPPW" >>slapd.conf
   ```

   The basic configuration of your LDAP server is complete and you are ready to start your LDAP server.
6. Use the **service** command at the Linux prompt to start the LDAP server:

   ```
   # service ldap start
   ```

   You should receive a green OK. If not, check for error messages in the /var/log/messages file that relate to the slapd and then run the command again.
7. Configure the LDAP server to start automatically on boot, using the **chkconfig** command:

   ```
   # chkconfig --level 235 ldap on
   ```
8. Make sure that the LDAP server is running and responding to queries, using the **ldapsearch** command:

   ```
   # ldapsearch -h localhost -x -b <base_suffix>
      '(object class=*)'
   ```

   No entries should be returned, though you can expect a positive response from the LDAP server:

   ```
   # ldapsearch -h localhost -x -b o=yourOrg '(objectclass=*)'
   version: 2
   # filter: (objectclass=*)
   # requesting: ALL
   # search result
   search: 2
   result: 32 No such object
   # numResponses: 1#
   ```

   If the LDAP server responded correctly to the query, you are ready to configure your LDAP server to work with SAN File System.

Continue with *Configure OpenLDAP for SAN File System*.

## Configure OpenLDAP with SAN File System

This topic lists the steps required to configure OpenLDAP to be used with SAN File System.

Complete *Install OpenLDAP packages*, *Configure the OpenLDAP client*, and *Configure the OpenLDAP server* before starting this procedure.

This procedure uses an example with a base suffix of “o=yourOrg” and root DN of “cn=Manager,o=yourOrg”.

1. Enter the *stdin* input mode of the **ldapadd** command:

```
# ldapadd -x -W -h localhost -D "cn=Manager,o=yourOrg"
Enter LDAP Password:       (<----- INPUT PASSWORD HERE )
```

2. Enter your root DN password as prompted. This is the password that you entered at step 3 of "Configuring OpenLDAP server". If you entered your password correctly, you do not see a prompt. This indicates that the **ldapadd** command is waiting for you to type input at the keyboard.
3. Add the entry for the base suffix while in this mode:

```
# ldapadd -x -W -h localhost -D "cn=Manager,o=yourOrg"
Enter LDAP Password: (<----- INPUT PASSWORD HERE )
dn: o=yourOrg
objectClass: organization
o: yourOrg (<=== 2ND ENTER)
adding new entry "o=yourOrg" (<----- PRESSED Ctrl+D)
#
```

4. When the base suffix has been entered, press **Enter** a second time to indicate the end of the entry.
5. Press Ctrl+D to exit from the input mode.
6. Use the **ldapsearch** command to verify that the entry was added to the LDAP database:

```
# ldapsearch -x -h localhost -x -b o=yourOrg '
(objectclass=organization)'
```

7. Import your LDAP configuration using the **ldapadd** command:

```
# sed "s/Example/yourOrg/" sfsExample.ldif > sfsbase.ldif
```

   **Note:** After you have changed the organization, you have to modify the userPassword field for each user as required in the sfsbase.ldif file.

8. Import your entries in the file with the **ldapadd** command:

```
# ldapadd -x -W -h localhost -D "cn=Manager,o=yourOrg"
-f sfsbase.ldif
Enter LDAP Password:
adding new entry "cn=Manager,o=yourOrg"
adding new entry "ou=Users,o=yourOrg"
adding new entry "cn=yourOrgAdmin Administrator,ou=Users,
o=yourOrg"
adding new entry "cn=yourOrgMon Monitor,ou=Users,o=yourOrg"
adding new entry "cn=yourOrgBack Backup,ou=Users,o=yourOrg"
adding new entry "cn=yourOrgOper Operator,ou=Users,
o=yourOrg"
adding new entry "ou=Roles,o=yourOrg"
adding new entry "cn=Administrator,ou=Roles,o=yourOrg"
adding new entry "cn=Monitor,ou=Roles,o=yourOrg"
adding new entry "cn=Backup,ou=Roles,o=yourOrg"
adding new entry "cn=Operator,ou=Roles,o=yourOrg"
```

9. Enter your root DN password when prompted, which is the same as you entered in step 3 in "Configuring OpenLDAP server".

   **Note:** Adding an entry twice fails at that point, and any subsequent entries are not processed. If some entries are correct and others failed, only attempt to add the objects that failed.

10. Use the **ldapsearch** command again to verify the objects. Refer to the example in step 6.
11. The LDAP directory (ldbm) files reside in the directory /var/lib/ldap/ by default. Verify that they exist by entering the following command:

```
# ls -lt /var/lib/ldap/
```

    The output appears similar to the following example:

```
total 56
-rw------- 1 ldap ldap 8192 Sep 21 16:41 cn.dbb
-rw------- 1 ldap ldap 8192 Sep 21 16:41 dn2id.dbb
-rw------- 1 ldap ldap 8192 Sep 21 16:41 id2entry.dbb
-rw------- 1 ldap ldap 8192 Sep 21 16:41 nextid.dbb
-rw------- 1 ldap ldap 8192 Sep 21 16:41 objectClass.dbb
-rw------- 1 ldap ldap 8192 Sep 21 16:41 sn.dbb
-rw------- 1 ldap ldap 8192 Sep 21 16:41 uid.dbb
```

**Tip:** If you want to reconfigure the LDAP directory from scratch, stop slapd, remove the ldbm files, start slapd, then begin the steps in this procedure.

OpenLDAP is now configured and ready to be used with SAN File System.

## Configuring LDAP using Microsoft Active Directory LDAP

These instructions illustrate a simple method of configuring Active Directory to be used as the LDAP architecture for SAN File System. You can either use these instructions to get Active Directory running to support SAN File System, or to understand one way to configure SAN File System to work with your Active Directory environment.

This topic assumes the use of Windows 2000 Server. Other Windows Server versions are similar but not identical. SAN File System accesses Active Directory using the standard LDAP protocol. It does not rely on Microsoft extensions to LDAP standards. For more information about Active Directory, refer to www.microsoft.com/ad. Also, you are responsible for making sure that your use of Microsoft Active Directory is covered by the appropriate license.

The overall steps required to configure Active Directory to be used with SAN File System are:

- Install Active Directory.
- Configure Active Directory with a domain containing groups and users.
- Configure SAN File System to access the Active Directory installation.
- Validate SAN File System access to Active Directory.

### Installing Active Directory

The steps in this section describe how to enable the Windows system to be a domain controller and define a domain within which SAN File System users and groups exist. This example creates a domain called sanfsdom.net to contain the directory elements needed to manage SAN File System.

You must have a Microsoft Windows 2000 Server system.

If you plan to use SAN File System with a preexisting Active Directory installation, you can skip this section.

1. Open the Server configuration program using **Start** → **Programs** → **Administrative** → **Configure your server**. The Windows 2000 Server system configuration is displayed.
2. Click **Active Directory** in the left menu. The Active Directory Installation wizard appears.
   a. Click **Next** to continue.
   b. Select the *Domain controller for a new domain* option and click **Next**.
   c. Select the *Create a new domain tree* option. In this simple example, a new domain is sufficient.

d. Select the *Create a new forest of domain trees* option for this example. Blending a SAN File System domain into other domains requires Active Directory expertise.
e. Specify the name of your new domain.
f. Specify the database and log locations.
g. Specify the folder for the shared system volume.
h. If you are using an Active Directory domain name that is not known to DNS, then the following message appears and is normal in this case. For a minimal configuration, DNS recognition of the Active Directory domain is not necessary. Click **OK**.

```
The wizard cannot contact the DNS server that handles the name
"sanfsdom.net" to determine if it supports dynamic update.  Confirm you
DNS configuration, or install and configure a DNS server on this computer.
```

i. The *Permissions* options on the next screen do not affect SAN File System. Click **Next**.
j. Enter a password twice.
k. Click **Finish** and follow the prompts to reboot.

Continue with *Configuring Active Directory*.

## Configuring Active Directory

This topic lists the procedures required to configure Active Directory for use with SAN File System.

You must complete the steps in *Installing Active Directory* before completing this procedure.

The following objects must be configured by the domain that is used by SAN File System:

- A user for the SAN File System Administrative agent to access the contents of the Active Directory instance. In the following example, this user is called LDAP_Admin.
- A Global Security Group representing the SAN File System Administrator role. In the following example, this group is called SANFS_Admins.
- A Global Security Group representing the Operator role (optional).
- A Global Security Group representing the Backup role (optional).
- A Global Security Group representing the Monitor role (optional).
- One or more users who are authorized in the Administrator role.
- Optionally, one or more users are authorized in the Operator role.
- Optionally, one or more users are authorized in the Backup role.
- Optionally, one or more users are authorized in the Monitor role.

1. Add these elements using the Active Directory Users and Computers interface.
   a. Open the Active Directory Users and Computers interface by clicking **Start → Programs → Administrative Tools → Active Directory Users and Computers**. The Active Directory Users and Computers interface, has two panels. The left panel shows the tree type and the right panel shows existing objects of the type highlighted in the left panel.
   b. In the left panel, click + next to the sanfsdom.net domain to expand it and show its types.
   c. Click **Users type** to show the users that exist within the domain.
2. Add the LDAP_Admin user

a. With the **Users type** highlighted (selected) in the left panel in the Active Directory Users and Computers interface, select **Action** → **New** → **UserTo**.
b. Add the LDAP_Admin user, and fill in the First name and User logon name and click **Next**. SAN File System uses the User logon name in its LDAP traversal of Active Directory.
c. Select the *Password never expires* option and enter the password. Do not select the *Disable account* or *User must change password* options. The Administrative agent automatically uses this user and password combination to access Active Directory, so future password changes must be made within Active Directory and the tank.properties file at the same time. (Note that the Administrative agent must be restarted any time a tank.properties value is changed as described in 5 on page 39.)
d. Click **Next**, then **Finish**. The newly created user appears in the object list.

3. Add the SAN File System administration group There are four SAN File System Administration groups, corresponding to the SAN File System administration roles: Administrator, Backup, Operator, and Monitor.
   a. With the **Users type** highlighted in the Active Directory Users and Computers interface, click **Action** → **New** → **Group**.
   b. Fill in the group name. It should be a Global Security group.
   c. Click **OK**.
   d. Modify the newly created group to specify its Description property. In the example configurations, the Description property is used by the Administrative agent in searching Active Directory, so it must be the verbatim string corresponding to the SAN File System role, in this case "Administrator" with no trailing spaces.
   e. Click **OK**.
   f. Repeat the steps in this section to create the groups for each of the other three SAN File System roles (Operator, Backup, and Monitor), in each case modifying the Description property to match the SAN File System role exactly. The other three roles are not necessary to enable basic SAN File System administration. If used, they provide restricted levels of capability within the SAN File System GUI and CLI.
4. Create users authorized to manage SAN File System. To create an authorized user, you must first create the user, and then specify that it is a member of one of the SAN File System administration groups created in the previous section.
   a. To create a user, follow the same steps described in "Adding the LDAP_ Admin user", substituting the user login name that you want to use into the First name and User logon name fields. The password that you specify is the password that must be given to tanktool and the SAN File System console for authentication. If you use the **tankpasswd** command to specify an administrator password on the SAN File System cluster, it needs to be changed to match the password specified for the authorized user in Active Directory.
   b. After creating the new users, you can create membership in one of the four SAN File System administration groups using one of the following methods:
      - Double-click the group in the Active Directory Users and Computers interface, select the Members tab in the group properties panel, select the user that you want to authorize, and click **Add**.
      - Double-click the user, select the Member-Of tab in the user properties panel, select the group in which you want to include the user, and click **Add**. Then click **OK**.

Continue with *Configuring SAN File System to use Active Directory*.

## Configuring SAN File System to use Active Directory

This topic lists the modifications to SAN File System that are necessary to enable it to use Active Directory as its LDAP architecture.

You must complete the steps in *Installing Active Directory* and *Configuring Active Directory* before completing this procedure.

1. You have to modify the tank.properties configuration file on each SAN File System engine to reflect the Active Directory-compatible LDAP configuration. You can modify this file at SAN File System initialization using a special setupTank flag, or later by editing the tank.properties file directly.

   **Attention:** Running setupTank or modifying the tank.properties file incorrectly can render an existing SAN File System installation unusable.
2. To modify the SAN File System LDAP configuration during initialization, use the **-debug** parameter on the **setupTank** command to access and modify the tank.properties file as part of the setup script:

   ```
   setupTank -debug -setmaster
   ```
3. If you are not modifying the LDAP during configuration, use a text editor such as vi to modify the tank.properties file on each engine:

   ```
   vi /usr/tank/admin/config/tank.properties
   ```
4. Make changes as required to the following variables in the tank. properties file:

   **LDAP_SERVER=IP-address**
   For IP-address, specify the IP address of the Windows 2000 system running the Active Directory instance.

   **LDAP_USER=cn=LDAP_Admin,cn=Users,dc=sanfsdom, dc=net**
   For LDAP_Admin, specify the User logon name of the user created for the Administrative agent in searching Active Directory. For sanfsdom and net, substitute the parts of the domain name chosen for your SAN File System users and groups. The example uses sanfsdom.net. Represent more dotted-domain qualifications with more dc= clauses. The cn=Users clause represents the object type of this object. It is recommended that you do not change the object type unless you have Active Directory expertise.

   **LDAP_PASSWD= password**
   For password, substitute the password given to the LDAP_Admin user.

   **LDAP_BASEDN_ROLES=cn=Users, dc=sanfsdom,dc=net**
   This variable identifies the user object enabling the administrative agent to search Active Directory. Modify dc=sanfsdom,dc=net to indicate the domain name chosen for your SAN File System users and groups. The cn=Users clause represents the object type of this object. It is recommended that you do not change the object type unless you have Active Directory expertise.

   **LDAP_ROLE_MEM_ID_ATTR=member**
   This variable indicates to SAN File System the name of the attribute that relates the object representing a SAN File System administration role (for example, SANFS_Admins) to the user objects authorized for that role. When you use Active Directory with its default schema, this attribute must be “member”.

   **LDAP_USER_ID_ATTR=sAMAccountName**
   This variable indicates to SAN File System the user object attribute that

contains the logon name. This name corresponds to the field Active Directory calls User logon name in the New Object and Object Properties panels for User objects. It is recommended that you do not change the attribute type unless you have Active Directory expertise.

**LDAP_ROLE_ID_ATTR=description**
This variable indicates to SAN File System the group object attribute that contains the SAN File System role name. This name corresponds to the field AD calls Description in the New Object and Object Properties panels for Group objects. It is recommended that you do not change the attribute type unless you have Active Directory expertise.

**LDAP_SECURED_CONNECTION=false**
This variable indicates to SAN File System not to use SSL to connect to Active Directory.

**LDAP_CERT=**
This variable indicates to SAN File System information about the certificate needed by the AD instance to establish an SSL connection.

**Note:** The user name and passwords used by administrators to access sfscli and the SAN File System Console must match the ones specified in a user object in the member relation to one of the group objects representing administration roles. See the IBM TotalStorage SAN File System Administrator's Guide and Reference, (GA27-4317) for information about the methods for specifying passwords to these interfaces. If you are performing an initial installation of SAN File System, the administrator user name and password can also be set in a dot file (.tank.passwd) used by sfscli in response to prompts from within setupTank.

5. After you change the tank.properties file, the administrative agent must be restarted on each engine in the cluster for the new settings to take effect. Use the **stopCimom** and **startCimom** commands to restart the administrative agent.

   **Note:** Using this simplified method for configuring Active Directory and SAN File System results in warning messages in the administrative agent log. The administrative agent generates the warning "Role name xxx is invalid" for each user object that exists in the domain that does not have a description matching one of the SAN File System role names. For this reason, following the example in this paper is not practical for a domain that is used for other applications besides SAN File System. The method presented here has to be adapted to be useful in larger scale Active Directory domains.

Continue with *Validating the Active Directory and SAN File System configurations*.

## Validating the Active Directory and SAN File System configurations

This topic explains how to validate that the Active Directory and SAN File System configurations are set correctly.

You must complete the steps in *Installing Active Directory*, *Configuring Active Directory*, and *Configuring SAN File System to use Active Directory* before completing this procedure.

In the SAN File System administrative CLI (tanktool), the **lsadmuser** command lists the contents of Active Directory, relevant to SAN File System, found by the Administrative agent. Compare this listing to the roles and authorized users that you have entered.

1. Run the following command:

   ```
   ldapsearch –h ldap://LDAP_SERVER/ –w 'password"
   –D "LDAP_USER" -x –b "LDAP_BASEDN_ROLES" '(objectclass=group)'
   ```

   On the SAN File System engine, this command uses the LDAP_USER login as described earlier to list all group objects on the domain server, LDAP_ SERVER, that match a LDAP_BASEDN_ROLES object. Based on our example, the command might be:

   ```
   ldapsearch -h ldap://128.47.79.140/ -w "password"
   -D "cn=LDAP_Admin,cn=Users,dc=sanfsdom,dc=net"
   -x -b "cn=Users,dc=sanfsdom,dc=net" '(objectclass=group)'
   ```

   Here is an example of the output of the command with the details of other groups removed:

   ```
   CN=SANFS_Admins,CN=Users,DC=sanfsdom,DC=net
   member=CN=newuser1,CN=Users,DC=sanfsdom,DC=net
   member=CN=IFTEST_USER,CN=Users,DC=sanfsdom,DC=net
   member=CN=stuser,CN=Users,DC=sanfsdom,DC=net
   member=CN=root,CN=Users,DC=sanfsdom,DC=net
   info=This global security group designates users who have SANFS
   Administrator authorization. cn=SANFS_Admins
   description=Administrator
   groupType=-2147483646
   instanceType=4
   distinguishedName=CN=SANFS_Admins,CN=Users,DC=sanfsdom,DC=net
   objectCategory=CN=Group,CN=Schema,CN=Configuration,
   DC=sanfsdom,DC=net
   objectClass=top
   objectClass=group
   objectGUID=NOT ASCII
   objectSid=NOT ASCII
   name=SANFS_Admins
   sAMAccountName=SANFS_Admins
   sAMAccountType=268435456
   uSNChanged=2756
   uSNCreated=2744
   whenChanged=20031106013743.0Z
   whenCreated=20031106005502.0Z
   CN=SANFS_Operators,CN=Users,DC=sanfsdom,DC=net
   cn=SANFS_Operators
   groupType=-2147483646
   instanceType=4
   distinguishedName=CN=SANFS_Operators,CN=Users,DC=sanfsdom,D
   C=net
   objectCategory=CN=Group,CN=Schema,CN=Configuration,DC=sanfs
   dom,DC=net
   objectClass=top
   objectClass=group
   objectGUID=NOT ASCII
   objectSid=NOT ASCII
   name=SANFS_Operators
   sAMAccountName=SANFS_Operators
   sAMAccountType=268435456
   uSNChanged=2787
   uSNCreated=2785
   whenChanged=20031106145326.0Z
   whenCreated=20031106145326.0Z
   ```

2. The '(objectclass=group)' suffix limits the output to Group objects only. Note the member relations shown for the SANFS_Admins group. The entry for SANFS_Operators, by contrast, does not show any members added. A similar command can be run to see the users in the Active Directory domain. The **ldapsearch** command without any parameters prints a usage statement describing all the command options.

Active Directory is now set up and working correctly with SAN File System.

## Preparing the engine for installation

This topic describes how to prepare a metadata server engine for installation of software. You need to complete these steps for each metadata server.

1. Ensure that the engine, including the RSA II adapter is cabled properly. See "Cabling" on page 42.

   **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."
2. Obtain software to be used on the installation process. See "Obtain prerequisite software" on page 46.
3. Update the system BIOS, if necessary.
4. Use the LSI Logic Configuration Program documentation provided with the engine to mirror the boot drive. (This step applies only when you are using an LSI SCSI RAID controller.)

### 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**.

## Cabling

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

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 8 on page 44 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.

      **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 8 on page 44 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 8 on page 44.
      3) Connect an 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 6. and Figure 7 on page 43.
      4) The 9-pin D-shell serial connector on the ASM connector is not used.

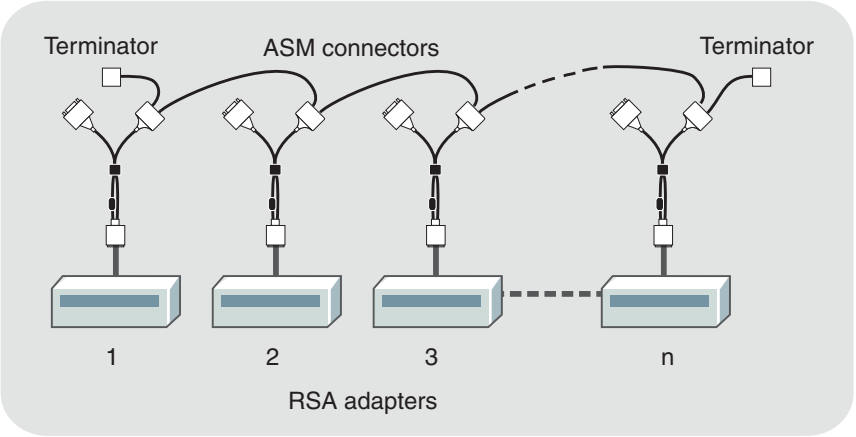


*Figure 6. Connecting the RSA adapters together*

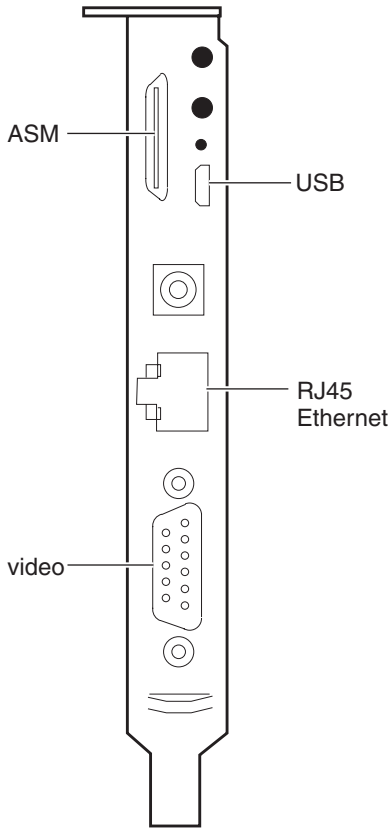


*Figure 7. RSA II adapter connectors*

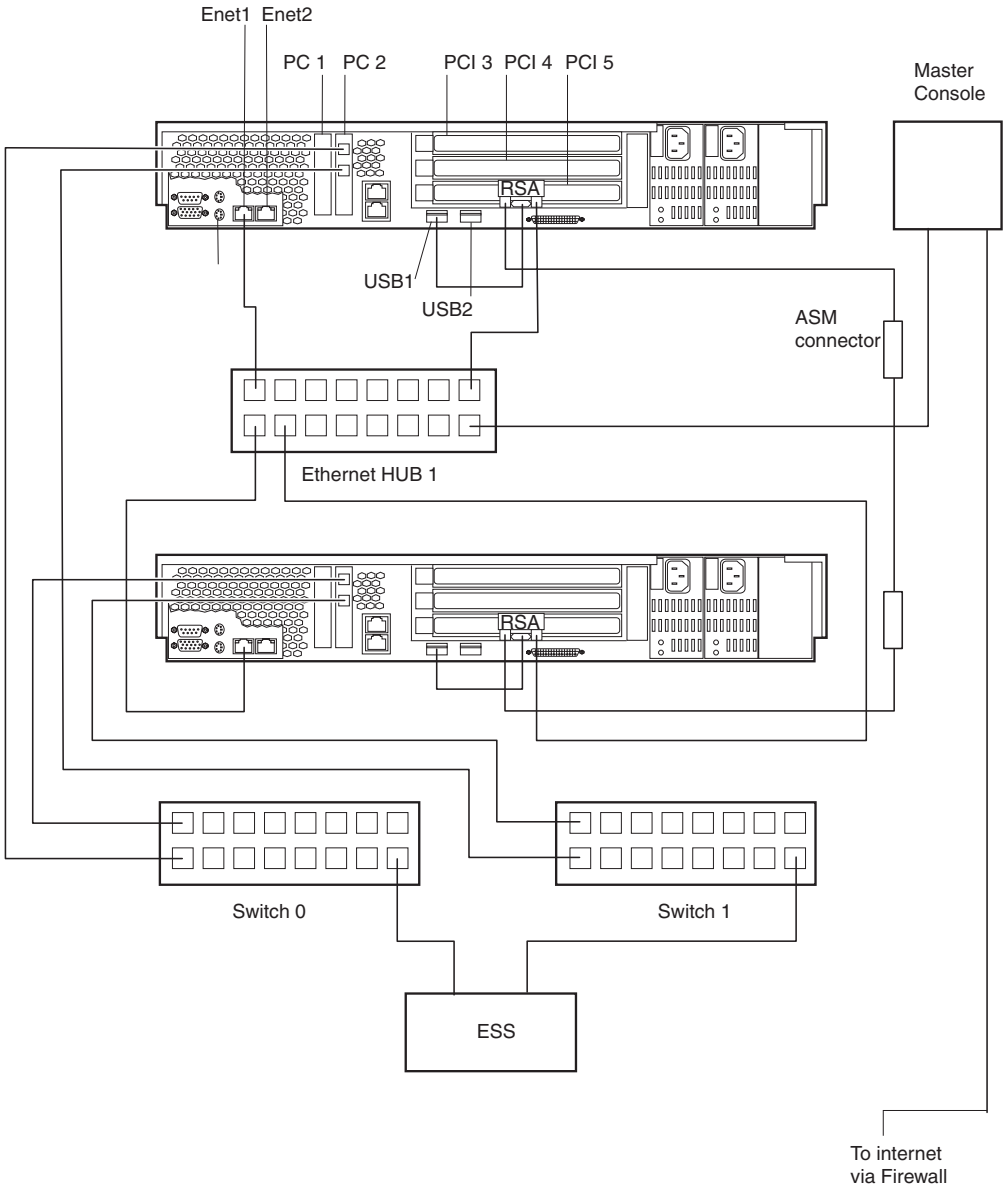


*Figure 8. Two-node, two-switch, two-hub cabling example*

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.

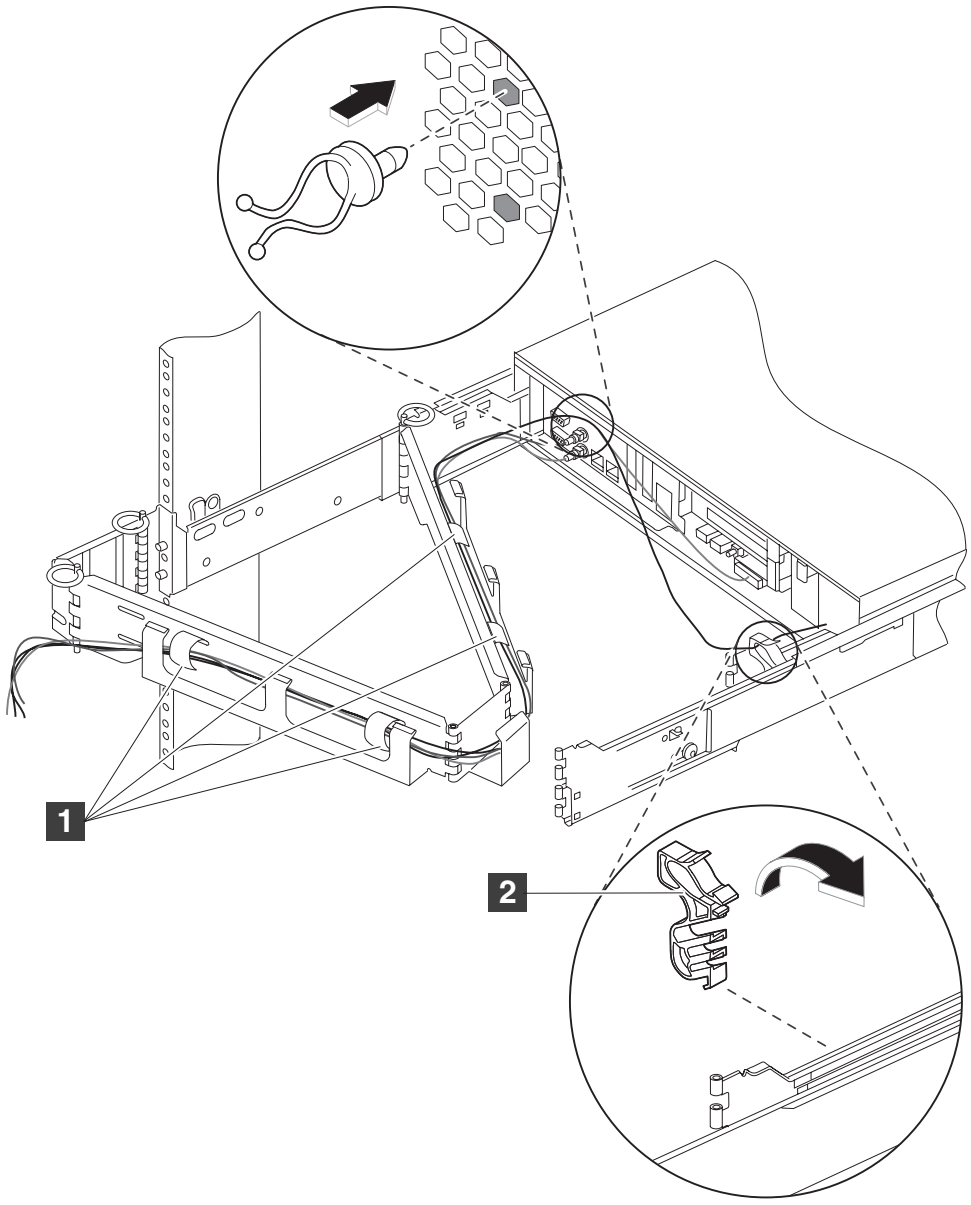


*Figure 9. Attaching the cable straps*

**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 8 on page 44 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 need is available on the SAN File System CD-ROM.

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

- SUSE Linux Enterprise Server 8. You need a licensed copy of SUSE Linux Enterprise Server 8 for each of the metadata server engines in the cluster. For more information about obtaining SUSE Linux Enterprise Server 8, see 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 1.17 (GEJ57B). To determine the current version of the system BIOS, reboot the engine and watch for the BIOS version to be displayed. Note that the BIOS version level on the IBM support Web site, 1.17, is equivalent to BIOS level 1.1.7, as shown in the example in "tmvt" on page 136.

To obtain the correct version of the system BIOS, see 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 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 1.17 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 need a diskette for this firmware update as well.

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

```
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.

# Chapter 3. Metadata server engine setup

This topic provides procedures for setting up the metadata server engine. Note that you need to complete these procedures on each metadata server.

Use the following checklist to set up the metadata server engine.

**Note:** This topic describes the procedures for setting up the metadata server engine on the xSeries 345. For details on the differences between setting up the metadata server engine on the xSeries 345 and xSeries 346, see "xSeries 346 installation upgrade procedure" on page 54.

| Steps | | For more information... |
|---|---|---|
| 1 | Set up the metadata server. | "Setting up the metadata server" |
| 2 | Install the operating system. | |
| a | Install SUSE Linux Enterprise Server 8. | "Installing the operating system" on page 48 |
| b | Disable the X Window System. | "Disabling the automatic starting of the X Window System" on page 51 |
| c | Set the date and time on the engine. | "Setting the time and date on the Metadata servers" on page 52 |
| d | Apply United Linux Service Pack 3 (SP3) updates. | "Apply United Linux Service Pack 3 (SP3) updates" on page 53 |

## Setting up the metadata server

This topic provides an overview of the steps required to set up the SAN File System metadata server. You need to complete these steps on each metadata server.

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

- You have unpacked the engine and installed (but not configured) the RSA II.
- 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 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 41.
2. Install all software, including the operating system, prerequisite software, and the SAN File System software. See "Installing the operating system for the metadata server" on page 48.

## Installing the operating system for the metadata server

This topic provides an overview of the steps that are required to install the operating system for the metadata server. You need to complete these steps on each metadata server.

1. Install SUSE Linux Enterprise Server 8. See "Installing the operating system"
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 51
3. Set the time and date on the metadata server engine. See "Setting the time and date on the Metadata servers" on page 52
4. Apply the United Linux Service Pack 3. See"Apply United Linux Service Pack 3 (SP3) updates" on page 53.

### Installing the operating system

This topic describes the procedure for installing SUSE Linux Enterprise Server 8 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 53 to apply the Service Pack. Do not use the United Linux Service Pack 3 CD Kickstart functionality because the installation may not complete successfully.

The SUSE Linux package contains three sets of CDs:
- SUSE Linux Enterprise Server (CD1)
- United Linux for SUSE Linux Enterprise Server 8 (CD1, CD2, CD3)
- SUSE Linux Enterprise Server SP (CD1, CD2)

**Tip:** Many of parameters that you configure during the SUSE installation (network interface) can be viewed or changed following the installation by entering **yast** at the command line. The YaST interface is similar to the smit interface provided in UNIX.

1. Disconnect all SAN (fibre channel) cables from the metadata server.

   **Attention:** Failure to perform this step can result in loss of SAN File System metadata.
2. Insert the SUSE Linux Enterprise Server SP CD 1 into the CD-ROM drive and reboot the engine.
3. When the install program displays *Make sure CD 1 is in the drive!*, eject the SUSE Linux Enterprise Server SP CD 1 and insert the standard SLES 8 CD 1.
4. When you see the installation 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.
5. Make sure that Installation is selected.
6. Enter the following kernel options in the boot-options text field:

   ```
   acpi=oldboot vga=normal
   ```

7. Using the resolution function keys, select the screen resolution that matches your monitor.
8. Press **Enter** to continue the installation.
9. Read the SUSE End User License For SLES agreement and click **Accept**.
10. Select the language and click **Accept**.
11. When the type of installation pop-up window appears, select **New installation** and click **OK**.
12. Click the **Change...** dropdown menu, and then select **Software** from the dropdown list.
13. Change to **Default system for UnitedLinux**. If you are prompted to confirm that you really want to reset your detailed selection, click **Yes**.
14. Click **Detailed selection...**
15. From the selection panel, click **C/C++ compiler and tools** to add this package to the software list. Then click **Accept**.
16. Click the **Change...** dropdown menu again, and then 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.

| Device | Size | Type | 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.

17. Click **Change...** to modify other settings, such as the timezone. After modifying these settings click **Accept**.
18. Click **Accept** to continue the installation.
19. At the warning prompt, click **Yes, install** to continue.
20. When prompted, insert the UnitedLinux CD-ROM 1 and click **OK**.
21. When prompted, insert UnitedLinux CD-ROM 2 in the CD-ROM drive and click **OK**.
22. When prompted, reinsert the SUSE LINUX Enterprise Server CD1 and click **OK**.
23. When prompted that the base system is installed, remove the CD-ROM and click **OK**.
24. During the SUSE install, the local root user is automatically created. When prompted, enter a new password for root and click **Next**.
25. You are prompted to create more local users. Create at least one additional local user to prevent errors during the installation. For example, consider creating a user called **guest** if you need only one additional user. Click **Next** after you have created any additional users.
26. You are prompted to either enable or disable the 3D graphics-capable card. Click **No** to disable the 3D graphics-capable card.
27. 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.
28. When prompted about detecting printers, select **Skip detection**.
29. Click the **Network Interfaces** link to configure the interfaces.

    The Network Interface panel consists of two lists: one at the top and one at the bottom. The bottom list displays the interfaces that are configured, while the top list displays the non-configured interfaces. During the SUSE installation, the system automatically configures the first primary interface for you (ent0), using the first interface it discovers, and thus the interface appears in the bottom list.

    If this automatic configuration (pairing between the primary interface and the NIC) is not acceptable for your environment, delete the configuration to place the NIC back in the non-configured interface list.

If you make an error in specifying the IP address or the IP address has changed after completing the SUSE installation, you can access the IP configuration panel through the YaST interface, by entering **yast** at the command prompt.

**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.

30. 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**.
31. 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 set text mode as the default desktop setting when you installed SUSE Linux Enterprise Server 8 (in step 27 on page 50), you must disable the X Window System 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
   ```

### Configuring NTP on the metadata server

This topic describes how to configure NTP on the metadata server.

Stop all applications that are currently running on the SAN File System client. To stop applications other than SAN File System applications, refer to the documentation that comes with the applications.

1. From the command line, enter:

   ```
   rpm -qa | grep ntp
   ```

   xntp-4.1.1-289 is displayed.
2. Display the NTP run levels. These are usually set to *off*.

   ```
   chkconfig --list xntpd
   ```
3. Stop the xntpd daemon when you have configured it to start when the server starts or have started it manually. Otherwise continue with the next step.

   ```
   /etc/init.d/xntpd stop
   ```
4. Activate the xntpd daemon on boot:

   ```
   chkconfig xntpd on
   ```
5. Display run levels. Run levels 2, 3, 5 are listed as *on*:

   ```
   chkconfig --list xntpd
   ```
6. Open the NTP file with the editor:

   ```
   vi /etc/ntp.conf
   ```
7. Scroll down about halfway through the file and after the line that reads:

   ```
   NTP server IP_address_of_NTP_server
   ```

   add the following line:

   ```
    server xx.xx.xx.xx #IP_address_of_server
   ```
8. If the ntp.conf file contains a line that reads:

   ```
   restrict default ignore
   ```

comment out this line. This line prevents the daemon from synchronizing with any of the servers.

9. Set your local time equal to the time on one or more servers:
   ```
   /usr/sbin/ntpdate -b <server_1> [<server_n>]
   ```
10. Set the hardware clock to the system time:
    ```
    /sbin/hwclock --systohc
    ```
11. Verify that the CMOS clock is correct:
    ```
    /sbin/hwclock --show
    ```
12. Start the xntpd daemon:
    ```
    /etc/init.d/xntpd start
    ```
13. Verify that the system date and time on the metadata server match the date and time on the NTP server:
    ```
    date
    ```

## 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 completing an initial installation of SUSE Linux Enterprise Server 8 on the metadata server engine, as described in "Installing the operating system" on page 48. 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.
   ```
   shutdown -r now
   ```
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
   ```

## Upgrading the Linux kernel

This topic describes the procedure for upgrading the Linux kernel on a metadata server.

Obtain the following packages through your SUSE SLES8 Maintenance Web service (or through a public Linux download site such as rpmfind.net):

- kernel-source-2.4.21-231.i586.rpm

- k_smp-2.4.21-231.i586.rpm

1. Download the packages to a directory of your choice, for example, /tmp.
2. From the directory to which you downloaded the packages, install the two packages with the following commands (ignore any “cannot remove directory” warnings):

   ```
   rpm -Uvh kernel-source-2.4.21-231.i586.rpm

   rpm -Uvh k_smp-2.4.21-231.i586.rpm
   ```

3. If you are using the lilo boot loader, be sure to edit the /etc/lilo.conf file to point to this new kernel, then run the **lilo** command.

   If you followed the standard SLES8 install instructions, you are using the grub boot loader.
4. Reboot the system
5. The new kernel reports itself as version 2.4.21-231. For example:

   ```
   evt5-mds2:/tmp  uname -a
   Linux evt5-mds2 2.4.21-231-smp 1 SMP Mon Jun 28 15:31:39 UTC 2004 i686 unknown
   ```

## xSeries 346 installation upgrade procedure

This topic presents details on installing the xSeries 346.

The xSeries 346 is functionally equivalent to the xSeries 345, but there are differences in physical layout and operating-system installation, as described in this topic.

**RSA II connectivity**

The connectivity of the xSeries 346 differs from the xSeries 345 in the following ways, and as shown in Figure 10 on page 55:

- The xSeries 346 uses an RSA II Slimline adapter that is not on the PCI bus, but is instead attached to an internal system connector. As such, there is no card edge for the RSA II, as there is in the xSeries 345. Instead, the RSA II Ethernet and ASM connections are positioned on the back panel of the xSeries 346.
- The external USB connection is not required for the xSeries 346, as that connection is made internally.
- There is no VGA connector specific to the RSA II, as there is on the xSeries 345. Use the regular video port on the left side of the back panel.
- The ASM ports are located in the center of the back panel of the xSeries 346 (stacked vertically). You can use these two ASM ports like the ports on the breakout cable (dongle) of the xSeries 345 RSA II. To do this, daisy chain the ports together with other metadata servers as described in Figure 6 on page 42.
- The RSA II Ethernet is located directly to the right of the regular system Ethernet ports.
- All other functions and configuration steps of the RSA II Slimline are the same as those for the regular RSA II.

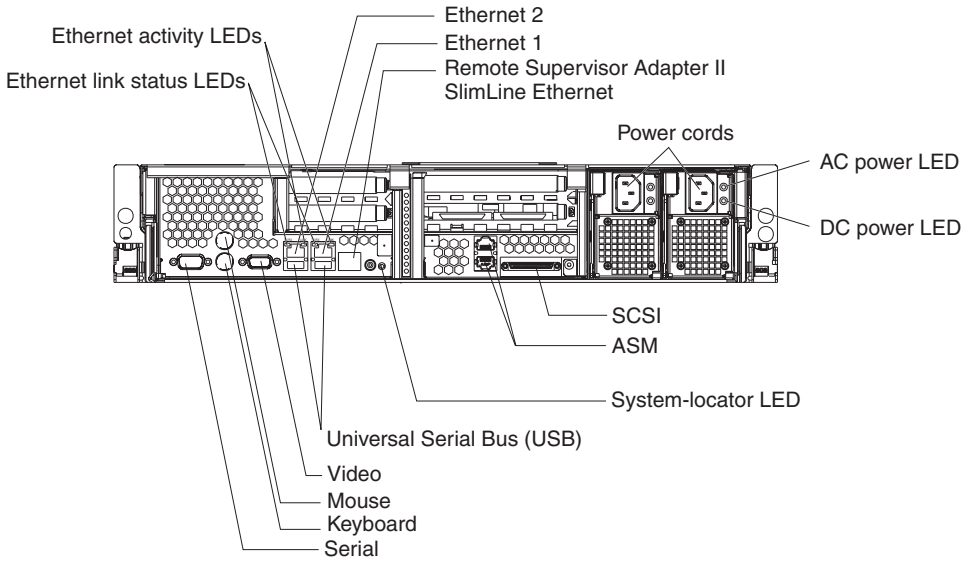


*Figure 10. Back plane of the xSeries 346*

**Onboard Ethernet adapters are Broadcom and require changes to SLES 8 installation**

The onboard network adapters for the xSeries 346 are Broadcom NetXtreme Gigabit, instead of the Intel™ E1000 in the xSeries 345. These adapters require some new Linux drivers that are not a part of the default SLES 8 Service Pack 3 installation. When installing the xSeries 346, take note of the following changes to the xSeries 345 installation steps presented in "Installing the operating system" on page 48:

1. Because networking is not available immediately after SLES 8 installation, you need to put the required drivers and rpms onto a CD. Before starting the metadata server installation procedure, obtain the following files and burn them onto a CD. In this procedure, this CD is referred to as the *Driver/Kernel CD*:
   - From SUSE:

     ```
     kernel-source-2.4.21-231.i586.rpm
     k_smp-2.4.21-231.i586.rpm
     ```

   - From Broadcom:

     ```
     bcm5700-7.3.5-1.src.rpm
     ```

     The Broadcom Web site is located at:

     ```
     http://www.broadcom.com/drivers/driver-sla.php?driver=570x-Linux
     ```

     Note that the download file is a zip file. You must open that zip file and extract the rpm contained in Server/Linux/Driver/bcm5700-7.3.5-1.src.rpm.
2. Disconnect all SAN (Fibre Channel) cables from the metadata server.

   **Attention:** Failure to perform this step can result in loss of SAN File System metadata.
3. Insert the SUSE Linux Enterprise Server SP3 CD 1 into the CD-ROM drive and reboot the engine.
4. When the install program displays *Make sure CD 1 is in the drive!*, eject the SUSE Linux Enterprise Server SP3 CD 1 and insert the standard SLES 8 CD 1.

5. At this point, continue with step 3 in "Installing the operating system" on page 48 and continue through "Apply United Linux Service Pack 3 (SP3) updates" on page 53.
6. After applying the SP3 updates in "Apply United Linux Service Pack 3 (SP3) updates" on page 53, upgrade the SAN File System 2.2 required kernel by completing these steps:
   a. Insert the Driver/Kernel CD that was prepared earlier into the CD drive and enter the following commands (ignore any error messages about an inability to delete a directory):

      ```
      mount /media/cdrom
      cd /media/cdrom
      cp kernel-source-2.4.21-231.i586.rpm /tmp
      cp k_smp-2.4.21-231.i586.rpm /tmp
      cp bcm5700-7.3.5-1.src.rpm /tmp
      cd
      umount /media/cdrom
      rpm -e kernel-source-2.4.21-138
      rpm -ivh /tmp/kernel-source-2.4.21-231.i586.rpm
      rpm -Uvh /tmp/k_smp-2.4.21-231.i586.rpm
      depmod -a
      ```

   b. Remove the CD from the CD drive.
   c. Reboot.
7. You can now install the Broadcom driver by entering the following commands:

   ```
   cd /usr/src/linux
   make mrproper
   cp /boot/config-2.4.21-231-smp .config
   make oldconfig
   make dep
   rpm -ivh /tmp/bcm5700-7.3.5-1.src.rpm
   cd /usr/src/packages
   rpm -bb SPECS/bcm5700.spec
   rpm -ivh --force RPMS/i386/bcm5700-7.3.5-1.i386.rpm
   ```

8. Edit the /etc/modules.conf file and change these lines (near the beginning of the file)

   ```
   alias eth0 off
   alias eth1 off
   ```

   to

   ```
   alias eth0 bcm5700
   alias eth1 bcm5700
   ```

9. Set up the dependency descriptions for the loadable kernel modules by creating the makefile and updating the kernel:

   ```
   depmod -a
   ```

10. After you have completed the installation of the driver, reboot.

After the reboot, the networking is operational. At this point, continue with the remaining SAN File System setup steps starting with "Installing prerequisite software on the metadata server engine" on page 57.

**Note:** Repeat these steps for all xSeries 346 metadata servers.

# Chapter 4. Metadata server software installation and configuration

This topic provides the procedures for installing prerequisite software and configuring the metadata server. Note that you need to complete these procedures on each metadata server.

Use the following checklist to install prerequisite software and configure each metadata server.

**Tip:** When entering a file name at the command line (for instance, with the **cp** or **rpm** commands) you can use the Tab key to automatically complete the file name. Enter enough of the file name to differentiate between files, and then press the **Tab** key.

| Steps | | For more information... |
|---|---|---|
| 1 | Install prerequisite software | "Installing prerequisite software on the metadata server engine" |
| a | QLogic driver. | "Install QLogic driver" on page 58 |
| b | MPCLI. | "Install MPCLI" on page 60 |
| c | Java runtime environment. | "Install the Java Runtime Environment" on page 60 |
| e | Eclipse. | "Install Eclipse" on page 61 |
| f | Ibmusbasm. | "Install ibmusbasm" on page 61 |
| g | Openslp. | "Install OpenSLP" on page 61 |
| h | IBM SDD or RDAC. | "Install the IBM Subsystem Device Driver (SDD)" on page 61 or "Installing Redundant Disk Array Controller" on page 100 |
| i | WebSphere 5.0 Express. | "Install IBM WebSphere 5.0 Express" on page 61 |
| 2 | Install the SAN File System software. | "Installing SAN File System software" on page 62 |
| 3 | Configure the RSA II adapter. | "Configuring the RSA II" on page 63 |
| 4 | Upgrade RSA II firmware, if needed. | "Upgrading RSA II firmware" on page 65 |

## Installing prerequisite software on the metadata server engine

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" on page 58.

- Management Processor Command Line Interface (MPCLI). See "Install MPCLI" on page 60.
- IBM Java™ Runtime Environment. See "Install the Java Runtime Environment" on page 60
- Eclipse. See "Install Eclipse" on page 61.
- The ibmusbasm daemon. See "Install ibmusbasm" on page 61.
- Openslp. See "Install OpenSLP" on page 61.
- IBM Subsystem Device Driver (SDD). See "Install the IBM Subsystem Device Driver (SDD)" on page 61.
- IBM Websphere 5.0 Express. See "Install IBM WebSphere 5.0 Express" on page 61.

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 Web site and that you can use this driver with your storage subsystem. You must install the 2.4.21-231 kernel patch before starting this procedure; see "Upgrading the Linux kernel" on page 53 for information about installing the patch.

1. Download the QLogic driver package to the temporary directory and change to that directory. If you download the driver package onto a floppy disk, follow these steps to copy the package to the temporary directory:
   a. Change to the temporary directory:
      ```
      cd /tmp
      ```
   b. Mount the floppy disk:
      ```
      mount /media/floppy
      ```
   c. Copy the package file:
      ```
      cp /media/floppy/qla2x00-v6.06.64-dist.tgz /tmp
      ```
   d. Unmount the floppy disk:
      ```
      umount /media/floppy
      ```
2. Mount the SAN File System CD:
   ```
   mount /media/cdrom
   ```
3. Run the QLogic driver build script, mkqla.sh:
   ```
   /media/cdrom/mkqla.sh /tmp/qla2x00-v6.06.64-dist.tgz
   ```

   **Note:**
   - By default, mkqla.sh searches for the qla*.tgz file in /root. If you copy the driver package into /root, you can run mkqla.sh without the **file** parameter.
   - The mkqla.sh script runs the **mkinitrd** command. Therefore, if you are using the lilo boot loader, be sure to run the **lilo** command before rebooting your system.
4. Attach the SAN (fibre channel) cables to the metadata server.
5. Reboot the engine:
   ```
   shutdown -r now
   ```
6. 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
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 02
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 03
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 04
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 05
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 06
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi2 Channel: 00 Id: 01 Lun: 07
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 00
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 01
  Vendor: IBM      Model: 2145             Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 02
  Vendor: IBM      Model: 2145             Rev: 0000
```

```
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 03
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 04
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 05
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 06
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 00 Lun: 07
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 00
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 01
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 02
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 03
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 04
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 05
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
Host: scsi3 Channel: 00 Id: 01 Lun: 06
  Vendor: IBM      Model: 2145            Rev: 0000
  Type:   Direct-Access                    ANSI SCSI revision: 03
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 blackslash (\) 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 has been packaged so that it installs under a different directory tree. To avoid unexpected results, use YaST2 to remove the SUSE-supplied version.

## 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 is 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-1.5.1.1-13.i686.ul1.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 are 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
```

## Install heterogeneous security

This topic describes how to install the heterogeneous security feature.

To take advantage of heterogeneous security in SAN File System, you must install and configure Winbind and Heimdal on each metadata server. Instructions for installing and configuring Winbind and Heimdal are provided in the /usr/tank/samba/scripts/INSTALL text file, which is available after extracting the hetsec_prereqs.tar file. The hetsec_prereqs.tar file is located on the SAN File System Software CD.

To install Winbind and Heimdal, perform these steps on each metadata server:

1. Copy hetsec_prereqs.tar to the /tmp directory:

```
cp hetsec_prereqs.tar /tmp
```

2. Change to the /usr/tank directory:

```
cd /usr/tank
```

3. Extract the hetsec_prereqs.tar file:

```
tar xvf /tmp/hetsec_prereqs.tar
```

4. Follow instructions in the /usr/tank/hetsec_prereqs/INSTALL file to complete the installation.

# 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. Note that though the version string for the combined package might differ from the version strings of the individual packages, this does not cause any problems with the installation.

   /media/cdrom/sfs-package-*build_level*.i386.rpm
3. When prompted, type the number that corresponds to your preferred language and press **Enter**.
4. When prompted, press **Enter** to view the International Program License Agreement.
5. After reading and agreeing to the license (by pressing **Enter** to page forward and typing **99** and pressing **Enter** to page backwards), type **1** and press **Enter** to accept the license agreement and install the software.

6. Change to the packages subdirectory.
   ```
   cd /usr/tank/packages
   ```
7. Install the administrative package.
   ```
   rpm -ivh sfs.admin.linux-build_level.i386.rpm
   ```
8. Install the metadata server package:
   ```
   rpm -ivh sfs.server.linux-build_level.i386.rpm
   ```
9. Run the Target Machine Validation Tool (TMVT) to verify that your hardware and software prerequisites have been met.
   ```
   /usr/tank/server/bin/tmvt -r
   ```

   The tool creates a default report file, tmvt.report (or you can pipe the results to a filename of your choosing). Examine the report file, paying particular attention to areas flagged as noncompliant. Resolve those prerequisites, and then rerun the tool until all prerequisites are compliant.

   **Notes:**

   a. TMVT noncompliance does not strictly prevent the installation of the SAN File System. It identifies deviations from the recommended hardware and software platform.
   b. If the kernel level on which you are running is not supported by SAN File System, you should correct the kernel version and run tmvt again.
   c. For more information about tmvt, see "tmvt" on page 136.

10. After you have installed the SAN File System software, you need to set the PATH environment variable in the .profile file to include the following paths:
    /usr/tank/server/bin
    /usr/tank/admin/bin

    One way to add paths to the PATH environment variable is by entering the following commands:

    a. cd $HOME
    b. vi .profile
    c. Press the *i* key to enter insert mode.
    d. Add this line to the .profile file:
       ```
       export PATH=$PATH:/usr/tank/server/bin:/usr/tank/admin/bin
       ```
    e. Press the escape key to exit insert mode.
    f. Enter :wq! to save the file and exit vi.
    g. Verify the entry:
       ```
       cat .profile
       ```
    h. Exit, and then log back in to make the new path take affect.

## Configuring the RSA II

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

Use the documentation supplied with the RSA II 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 need to supply the configuration setting when you run the **setupsfs** command).

**Note:** You can use the **setupsfs** command to configure the RSA II 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.

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

**Note:** If you have multiple RSA IIs on the network at the same time, each requires a unique IP address. If you do not define a unique IP address for each RSA II, when you use the Web interface to configure the RSA II, you are not able to determine the RSA II that you are updating.

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

Number the engines and RSAs 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 (PASSW0RD - 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. This name must match the name of the metadata server. You 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 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) or higher. You can use the Web interface to the RSA II to determine the firmware that you are currently using.

If the firmware is not at version 1.06, obtain the upgrade for the RSA II firmware at the following site. Note that the IBM Remote Supervisor Adapter Main Application version 1.06 on the IBM Support Web site is the equivalent of IBM Remote Supervisor Adapter Main Application Revision 16:

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

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

# Chapter 5. Creating the master and subordinate metadata servers

This topic provides the procedures for creating the master and subordinate metadata servers with the setupsfs utility.

## Setting up the master metadata server

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.

At this point in the installation, all the metadata servers are equally installed and configured. You need to select one of them to be the master metadata server.

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 **lsvpcfg** command and make a note of the virtual paths (vpaths) that can be used for metadata disks. You need this information when you are prompted for a list of the metadata disks, as shown towards the end of Table 6 on page 68. Place an *r* at the beginning of the vpath name before adding it to the list. For example, if a vpath name is *vpatha*, the metadata disk is expressed as

   ```
   /dev/rvpatha
   ```
3. Run the SAN File System setup utility. You are prompted to enter the information described in Table 6 on page 68. Accept the defaults for any prompts you see on the screen that are not described in the table.
   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/setupsfs -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/setupsfs -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 destroys all metadata stored by the SAN File System.**

   **Tip:** If you are using Active Directory as your LDAP server, you need to run setupsfs using the **-debug** parameter. Additional prompts are displayed. Enter the appropriate information for the additional LDAP prompts, and

accept the defaults for any prompts you see on the screen that are not described in Table 6.

*Table 6. 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 | Enter the DN of a user authorized to read the LDAP database, such as the LDAP Administrator or another user in the LDAP directory. |
| 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 if you used the example LDIF file enter `ou=Roles,o=yourOrg` . |
| LDAP member attribute | The attribute that contains the role members. For example if you used the example LDIF file, enter `roleOccupant`. |
| LDAP SSL 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`.<br>**Note:** When you run setupsfs, this certificate is embedded in the truststore. |
| RSA user name | The user ID used to access the RSA II adapter. The default is USERID. |
| RSA password | The passed for the user ID. The default is PASSW0RD (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. For example if you used the example LDIF file without any edits, you would enter `Admin`. |
| CLI password | The password defined for the CLI user. |
| Truststore password | The password that you define to create the truststore and later access it when you copy the truststore from the master metadata server to a subordinate metadata server. |
| 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/rvpatha /dev/rvpathb /dev/rvpathc /dev/rvpathd`. Create this list with the vpaths that you saw when you ran the **lsvpcfg** command. |

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

## 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
   root@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 setupsfs on the subordinate metadata server, the IP address and server name of the master metadata server are 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 also contains the LDAP certificate.

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

## Setting up the subordinate metadata servers

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

You 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/setupsfs
   ```

   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**.

# Chapter 6. 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/setupsfs -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/sfscli lsserver
   ```

   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
   ```

# Chapter 7. 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 systems:

- Windows 2000 Server and Advanced Server. See "Installing SAN File System on a Windows client."
- 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 76.
- Linux Red Hat Advanced Server 3.0. See "Installing SAN File System on a Linux client" on page 80.
- Sun Solaris 9 (64-bit). See "Installing SAN File System on a Solaris client" on page 83.

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. 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 2.2 software for a Windows client" on page 74.
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 a Windows client" on page 74.
4. Validate the installation. See "Validating the installation of SAN File System on a Windows client" on page 75.

5. Optionally configure the SAN File System Windows client to automatically restart when the client is rebooted. See "Automate client restart on reboot" on page 75

**Notes:** Ensure there is sufficient temporary space and root privileges when installing on a SAN File System disk.

When running various types of software installation and upgrade packages on Windows-based platforms, the installation software might select the SAN File System volume as the location to create or store temporary files or folders related to installation. This occurs because the SAN File System is often the disk with the most available space. For the installation or upgrade to function correctly, make sure that the following requirements are satisfied:

- The default storage pool backing up this root must have sufficient storage assigned to it. The install package creates files and folders under the root of the SAN File System volume.
- Use the administrative console to grant the SAN File System client root privileges. Only privileged clients can create files and folders under the root directory of a SAN File System volume.

## Obtain version 2.2 software for a Windows client

This topic explains how to obtain the version 2.2 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 file using the SAN File System console:

1. Start the SAN File System console from the client by opening a browser and entering the console Web address:

   `https://metadata_server_IP_addr:7979/sfs`
2. Enter the user name and password that you used when you configured LDAP (see "LDAP configuration" on page 22).
3. Select **Download Client Software**.
4. Follow the prompts to save the executable to a temporary directory.

## Installing the SAN File System software on a Windows client

This topic describes how to install version 2.2 of the SAN File System on a Windows client.

- The client for Windows can be installed only on a Windows 2000 server, Windows advanced server, or Windows 2003 standard server. A minimum version 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 is 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 2.2 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.

## 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

## 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.

You must be running AIX at the following operating levels:

- AIX 5.1 (32-bit only) uniprocessor or multiprocessor with maintenance level 3. 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.
- AIX 5.3 (32-bit and 64-bit) with bos.mp 5.3.0.0.

1. Obtain the SAN File System client software. See "Obtain version 2.2 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 an AIX client" on page 77.
4. Validate the installation of the client software. See "Validating the installation of SAN File System on an AIX client" on page 80.

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 2.2 software for an AIX client

This topic explains how to obtain the version 2.2 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. If you are installing the client on AIX version 5.3, the package name is sfs.client.AIX53.

You can load any 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.

**Attention:** Renaming the AIX package to a new name while downloading the package might cause a problem when you run the command:

```
installp -ac -d . client_package_name
```

command. The *client_package_name* must be the same name that was used for the package when it was built. If this is not the original package name, then some of the installation files might not be found.

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 on an AIX client

This topic describes how to install version 2.2 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.
  - AIX 5.3 (32-bit and 64-bit) with bos.mp 5.3.0.0.
- 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 see this choice only when you are disabling asynchronous I/O on AIX 5.2 or 5.3. 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 -ac -d dir client_package_name
```

where *dir* is the directory where you stored the package.

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 are prompted to enter values for the client configuration, as shown in Table 7. In most cases, you can accept the defaults.

*Table 7. AIX client configuration prompts*

| Parameter | Default | Description |
|---|---|---|
| kernextname | /usr/tank/client/bin/stfs | The client setup utility loads the SAN File System driver as a kernel extension and creates the file system driver instance. You must specify the path to the location of the SAN File System kernel extension. |
| devices | pat=/dev/rhdisk* | The SAN File System client determines which disks to use as SAN File System user data volumes by searching a list of disks, called device candidates. The device candidate list can be specified as a pattern or directory: pat=<*pattern*>, dir=<*directory path*> |
| clientname | hostname | The name can be any string, but must be unique among all SAN File System clients. |
| server_ip | No default | The SAN File System client must connect to one of the metadata servers in the cluster. After the client establishes a connection to the server, the server notifies the client of any other servers in the cluster. |
| server_port | 1700 | The SAN File System client must connect to the client-server port on the metadata server. Running the command **sfscli statserver –netconfig <*server_name*>** on the server displays the client-server port. |

*Table 7. AIX client configuration prompts (continued)*

| Parameter | Default | Description |
|---|---|---|
| mount_point | /mnt/sanfs | The client setup utility mounts the SAN File System to a specified mount point (directory) and creates the file system image. If the specified mount point does not exist, it is created.<br><br>**Attention:** Do not enter the mount point or directory of any general directories that are used by the base operating system of the client. For example, /, /root, /var, /etc, /usr. Doing this might cause the client operating system to stop performing basic functions. If you do mount the client at a standard directory, call the IBM Support Center for assistance. |
| readonly | No | If you mount the SAN File System as read-only, data and metadata in the file system can be viewed, but not modified. |
| majornumber | 99 | A major number is required to register the SAN File System driver with the kernel. Change the default only if the default major number is already in use. |
| convertertype | ISO-8859-1 | The NLS converter provides the metadata server with data on how to convert strings from the SAN File System client into UNICODE. Refer to the International Components for UNICODE Web site noted in “Web sites” on page vii for a list of supported converters. |
| nettype | tcp | The transport protocol determines how the SAN File System client connects to the Metadata server. Specify either *tcp* or *udp*. |
| stfserror | log | All SAN File System client errors are logged to the system log of the client machine. There are some error conditions that might require additional measures, such as when an application exits and a subsequent hardware failure prevents data from being committed to disk.<br><br>For these types of error conditions, you can select the *freezefs* or *systemhalt* options. The *freezefs* option prevents the SAN File System from writing additional data to disk and halts communication with the Metadata servers. The *systemhalt* option forces the client system to abruptly shut down. Choose either log, *freezefs*, or *systemhalt*. |
| verbose | No | By default, the client setup utility runs quietly, suppressing informational messages generated by the commands. You can choose to display these messages by entering *Yes*. |

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 either of the two types of Linux clients, RHEL or SLES. These steps must be performed on each Linux client in the SAN File System.

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. Obtain the SAN File System client software. See "Obtain version 2.2 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 a Linux client" on page 81.
4. Validate the installation. See "Validating the installation of SAN File System on a Linux client" on page 82.

## Obtain version 2.2 software for a Linux client

This topic explains how to obtain the version 2.2 SAN File System software for a Linux client using the graphical interface or command line interface.

This release of the SAN File System supports two Linux distributions:

- Red Hat Enterprise Linux 3.0 (RHEL)
- SUSE Linux Enterprise Server 8 (SLES8)

There is one Linux client package for each distribution. Each package contains support for all kernels that SAN File System uses for corresponding distributions. The package names are:

- sfs.client.linux_RHEL2.2.1-*n*
- sfs.client.linux_SLES82.2.1-*n*

Supported Linux kernel versions for RHEL are:

- 2.4.21-15.0.3.ELhugemem
- 2.4.21-15.0.3.ELsmp

The supported Linux kernel version for SLES8 is 2.4.21-231-smp.

You can load these packages onto the client from the SAN File System package repository. The package repository is located on each metadata server engine. Use either the secure copy function (scp) or the SAN File System console to transfer the package from a metadata server engine.

Transfer the package using secure copy:

1. On the client, change directories to a temporary directory.
   ```
   cd /tmp
   ```
2. Access the master metadata server.
3. 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 -R
   ```

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 on a Linux client

This topic describes how to install version 2.2 of the SAN File System on a Linux client

You can install the SAN File System client on the Red Hat Enterprise Linux 3.0 (RHEL) or SUSE Linux Enterprise Server 8 (SLES8) distribution.

1. Navigate to the directory where the client installation package is located.
2. Install the client package.
   ```
   rpm -ihv sfs.client.linux_RHEL2.2.1-n
   ```
   or
   ```
   rpm -ihv sfs.client.linux_SLES82.2.1-n
   ```
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 are prompted to enter values for the client configuration, as shown in Table 8. In most cases, you can accept the defaults.

*Table 8. Linux client configuration prompts*

| Parameter | Default | Description |
|---|---|---|
| devices | pat=/dev/sd*[a-z]* | The SAN File System client determines which disks to use as SAN File System user data volumes by searching a list of disks, called device candidates. The device candidate list can be specified as a pattern or directory: pat=<*pattern*> dir=<*directory path*>. |
| clientname | hostname | The name can be any string, but must be unique among all SAN File System clients. |

*Table 8. Linux client configuration prompts (continued)*

| Parameter | Default | Description |
|---|---|---|
| server_ip | No default | The SAN File System client must connect to one of the metadata servers in the cluster. After the client establishes a connection to the server, the server notifies the client of any other servers in the cluster. |
| server_port | 1700 | The SAN File System client must connect to the client-server port on the metadata server. Running the command **sfscli statserver –netconfig <*server_name*>** on the server displays the client-server port. |
| mount_point | /mnt/sanfs | The client setup utility mounts the SAN File System to a specified mount point (directory) and creates the file system image. If the specified mount point does not exist, it is created.<br><br>**Attention:** Do not enter the mount point or directory of any general directories that are used by the base operating system of the client. For example, /, /root, /var, /etc, /usr. Doing this might cause the client operating system to stop performing basic functions. If you do mount the client at a standard directory, call the IBM Support Center for assistance. |
| readonly | No | If you mount the SAN File System as read-only, data and metadata in the file system can be viewed, but not modified. |
| convertertype | ISO-8859-1 | The NLS converter tells the metadata server how to convert strings from the SAN File System client into Unicode. Refer to the International Components for UNICODE Web site noted in "Web sites" on page vii for a list of supported converters. |
| nettype | tcp | The transport protocol determines how the SAN File System client connects to the metadata server. Specify either **tcp** or **udp**. |
| etc_mtab | Yes | By default, when the file system mount succeeds, the client setup utility adds an entry for the file system image to /etc/mtab. Enter **No** if you do not want an entry added to /etc/mtab. |
| always_empty | No | By default, the number of blocks reported as free blocks by **statfs()** is actually the number of blocks in partitions that are not assigned to a fileset. Some programs might mistakenly report that there is no free space left in partitions assigned to the fileset, when there is actually free space available.<br><br>This option forces **statfs()** to report the number of free blocks as being one less than the number of blocks in the file system. |
| verbose | No | By default, the client setup utility runs quietly, suppressing informational messages generated by the commands. You can display these messages by entering **Yes**. |

**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 2.2 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 on a Solaris client."
4. Validate the installation. See "Validating the installation of SAN File System on a Solaris client" on page 85.

## Obtain version 2.2 software for a Solaris client

This topic explains how to obtain the version 2.2 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 on a Solaris client

This topic describes how to install version 2.2 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
   ```
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 are prompted to enter values for the client configuration, as shown in Table 9. In most cases, you can accept the defaults.

*Table 9. Solaris client configuration prompts*

| Parameter | Default | Description |
|---|---|---|
| devices | pat=/dev/dsk/c*t*d*s2 | The SAN File System client determines which disks to use as SAN File System user data volumes by searching a list of disks, called device candidates. The device candidate list can be specified as a pattern or directory. pat=<*pattern*> dir=<*directory path*>. |
| convertertype | ISO-8859-1 | The NLS converter provides the metadata server with data on how to convert strings from the SAN File System client into Unicode. Refer to the International Components for UNICODE Web site noted in “Web sites” on page vii for a list of supported converters. |
| clientname | hostname | The name can be any string, but must be unique among all SAN File System clients. |
| server_ip | No default | The SAN File System client must connect to one of the metadata servers in the cluster. After the client establishes a connection to the server, the server notifies the client of any other servers in the cluster. |
| server_port | 1700 | The SAN File System client must connect to the client-server port on the metadata server. Running the command **sfscli statserver –netconfig <*server_name*>** on the server displays the client-server port. |
| mount_point | /mnt/sanfs | The client setup utility mounts the SAN File System to a specified mount point (directory) and creates the file system image. If the specified mount point does not exist, it is created.<br><br>**Attention:** Do not enter the mount point or directory of any general directories that are used by the base operating system of the client. For example, /, /root, /var, /etc, /usr. Doing this might cause the client operating system to stop performing basic functions. If you do mount the client at a standard directory, call the IBM Support Center for assistance. |

*Table 9. Solaris client configuration prompts (continued)*

| Parameter | Default | Description |
|---|---|---|
| readonly | No | If you mount the SAN File System as read-only, data and metadata in the file system can be viewed, but not modified. |

## 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.

# Chapter 8. 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."
- Configure filesets to specify which metadata server manages a particular fileset and how much space can be allocated to files within a fileset. See "Configuring filesets" on page 88.
- Create placement policies to specify how files are placed in storage pools. See "Placement policies" on page 91.
- Optionally, you can migrate existing data to be managed by the SAN File System. See "Migrating data" on page 98.

## 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. On the master metadata server, add the master console as an SNMP manager.

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

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

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

Your SAN File System now sends 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.

Complete 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 can 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 fail and an SNMP alert is sent. If a soft quota is specified, then the allocation is allowed to succeed and an SNMP alert is 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
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 lost+found#
```

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 Linux Advanced Server 3.0.

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) and the **Description** field (for example, “a fileset for only Linux files”). Then, select a metadata server from the drop-down list (for example, ST0).

b. Under **Attach Point**, fill in the **Directory Path** field (for example, sanfs) and the **Directory Name** field (for example, linux30). 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
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 lost+found#
```

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-x-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 only Solaris 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
dr-xr-xr-x 2 1000000 1000000 4096 July 3 10:08 lost+found#
```

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-x-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, 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**. 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**. In the Client Sessions panel, select a client, select **Grant Clients Root Privileges** from the drop-down list, and then click **Go**.
4. To define the fileset owner, follow these steps:
   a. Open Microsoft Windows Explorer, expand the SAN File System drive letter, and then select the fileset you just created (for example, win2k).
   b. Set the owner by right-clicking and selecting **Properties**.
   c. Click the **Security** tab.
   d. Click **Advanced** , and then click the **Owner** tab.
   e. Select an owner from the **Change owner to** list and click **OK**.
5. Set permissions by selecting the folder that contains 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.

# Placement policies

This topic describes how to create placement policies to control where the data is placed.

Most rules in SAN File System policies are about placement. After an administrator has made choices about the administrative characteristics of the user storage pools used by the cluster, those characteristics are exploited through placement rules. This is accomplished by defining a rule that states *if the following condition is met, set the file's storage pool to be X*.

Files are often assigned to policies based on file names.

Files are placed in storage pools only when they are created. Changing the rules that apply to a file's placement does not cause the file to be moved.

Before migrating data, you must prepare the cluster for the addition of new files. When a SAN File System cluster is installed, there is an existing policy set called DEFAULT_POLICY. This policy has no rules: therefore, any SAN File System file created in any fileset goes to the default storage pool. You can list this policy set,

get the rules, and you can activate it using either the administrative CLI or the SAN File System console. You cannot modify or delete it.

## File placement policy syntax

This topic describes the syntax conventions for file-placement rules.

You can create a file containing policy rules for placing newly created files. You can then use this rule file when creating a policy using the **mkpolicy** command from the administrative CLI. You can also edit the policy rules that you create using the SAN File System console.

**Important:**

1. Every policy file must start with VERSION 1.
2. A policy is not required to contain any rules, in which case it would be equivalent to the default policy.
3. The maximum size of a policy is 32 KB.

You can also add comments to the policy. All comments must start with /* and end with */ (for example, /* comment */).

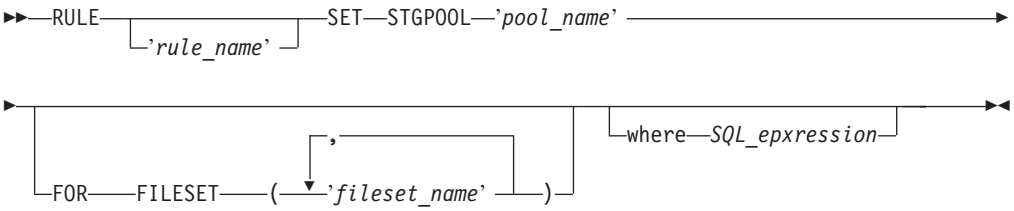


**Parameters**

**RULE**
Initiates the rule statement.

**'***rule_name***'**
Identifies the rule. This parameter is optional.

**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.

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

With SAN File System, you can use built-in functions that 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 percent (%) wildcard in the name to represent zero or more characters and use the underscore (_) wildcard to represent one single-byte or multibyte character.

**CREATION_DATE**
Date and time that the file was created.

**GROUP_ID**
Numeric group ID. This attribute is valid only for UNIX clients.

**USER_ID**
Numeric user ID. This attribute is valid only for UNIX clients.

**String functions**

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

**Important:** You must enclose strings in single-quotation marks. You can 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*. Both single-byte and multibyte characters count as one character in a string.

**CHAR_LENGTH(*x*)**
Determines the number of characters in string *x*. Both single-byte and multibyte characters count as one character in a string.

**CONAT(*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 UNIX-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 and the local time of the metadata server serving the directory within which the file is being created.

**Important:** Universal Time is used for all date and time functions.

**CURRENT DATE**
Determines the current date on the metadata server.

**CURRENT_DATE**
Determines the current date on the metadata server.

**CURRENT TIME**
Determines the current time on the metadata server.

**CURRENT_TIME**
Determines the current time on the metadata server.

**CURRENT TIMESTAMP**
Determines the current date and time on the metadata server.

**CURRENT_TIMESTAMP**
Determines the current date and time on the 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 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 can 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.

**Important:** 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 can 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 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'
```

# 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. For more information about the **migratedata** command, see "migratedata" on page 131.

**Attention:** When you migrate UNIX file system files to SAN File System, you lose any Posix-based access control lists (ACLs) from those files.

1. Estimate the time that it takes 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 data

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

For large data migrations, estimate the amount of time it takes 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 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.

Use the following steps to estimate data-migration time:

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. Enter 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 UNIX file system files to SAN File System, you lose any Posix-based 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. (The block size is displayed within the first several lines when the migration tool starts. The range is 1 Mb - 16 Mb.)

**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:** If you are not going to resume a particular data migration, you can clean up SAN File System after you stop that data migration by removing the migrated files (using standard utilities such as **rm** on UNIX or **del** on Windows) before you start a new data migration. However, if you intend to resume the migration, then do not remove the files that have been migrated.

## 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 is 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 receive an error and the migration stops. If you specify a different log file and specify the **–f** option, you receive a warning and the data on the target file system is 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.

If you deleted a set of data from the source, you can use the tool to put that set of data from SAN File System back to the source. But, the tool does not modify or delete anything in the source during migration.

# Installing Redundant Disk Array Controller

This procedure lists the steps required to install Redundant Disk Array Controller (RDAC). RDAC is needed only when you are using DS4000 (FAStT) devices.

1. Download the RDAC driver and Storage Manager 9.10 packages from either of the following Web sites:

   ```
   http://www-307.ibm.com/pc/support/site.wss/document.do?lndocid=MIGR-56707
   http://www-1.ibm.com/support/docview.wss?rs=572&uid=psg1MIGR-56707
   ```

   The packages, and their respective filenames, to download are:
   - IBM DS4000 (FAStT) Storage Manager for x86 Linux (26r0609.tgz)
   - IBM TotalStorage DS4000 (FAStT) Linux RDAC Software Package (rdac_linux_09.00.a5.00.tar.gz)
2. Move the Storage Manager package to a temporary directory and untar it:

a. Create the directory /tmp/rdac:

```
mkdir /tmp/rdac
```

b. Copy the 26r0609.tgz package into the /tmp/rdac directory:

```
cp 26r0609.tgz /tmp/rdac
```

c. Change to the /tmp/rdac directory:

```
cd /tmp/rdac
```

d. Untar the package:

```
tar -xvzf 26r0609.tg
```

3. Install the following Storage Manager RPMs:
   a. Change to the /tmp/rdac/Linux9p1/SM9ClientCode directory:

```
cd /tmp/rdac/Linux9p1/SM9ClientCode
```

   b. Install the run-time RPM:

```
rpm -i SMruntime-LINUX-9.10.A5.03-1.i586.rpm
```

   c. Install the utilities RPM:

```
rpm -i SMutil-LINUX-9.10.A5.01-1.i386.rpm
```

4. Before you unpack and install the RDAC driver, ensure that Auto Volume Transfer (AVT) is disabled and that you have the FW level of DS4000 (FAStT). Instructions and information are available in the RDAC readme file.
5. Set up the SUSE distribution:
   a. Install the kernel-source from the SUSE distribution:

```
rpm -U --force /tmp/server/kernel-source-2.4.21-231.i586.rpm
```

   b. Create a soft link to the kernel source:

```
ln -sf /usr/src/<linux-version>/usr/src/linux
```

   c. Ensure kernel version synchronization between the driver and the running kernel by entering the following commands in the Linux console window:
      1) Change to the linux directory:

```
cd /usr/src/linux
```

      2) Completely clean the kernel tree:

```
make mrproper
```

      3) Copy the new config file:

```
cp /boot/config-2.4.21-231-smp  .config
```

      4) Update the configuration using .config:

```
make oldconfig
```

      5) Rebuild the dependencies:

```
make dep
```

6. Untar the RDAC driver package:
   a. Change to the /tmp/rdac/ directory:

```
cd /tmp/rdac
```

   b. Untar the RDAC packages:

```
tar -zxvf rdac_linux_09.00.a5.00.tar.gz
```

7. Prepare the driver:
   a. Change to the /tmp/rdac/linuxrdac directory:

```
cd linuxrdac
```

   b. Remove the old driver modules in the linuxrdac directory:

```
make clean
```

c. Create a temporary file, SUSE-release. (The RDAC driver **make** command searches for this file to see what Linux distribution you are installing. This file is not automatically created when you install United Linux, so you need to temporarily create it.)

```
touch /etc/SUSE-release
```

d. Compile all driver modules and utilities:

```
make
```

8. Install the RDAC driver:
   a. Copy driver modules to the kernel module tree and build the new RAMdisk image (mpp.img) that includes RDAC driver modules and all driver modules that are needed during boot time:

```
make install
```

   b. Set up the dependency descriptions for the loadable kernel modules by creating the makefile and updating the kernel:

```
depmod -a
```

   c. Remove the temporary SUSE-release file:

```
rm /etc/SUSE-release
```

   d. Update grub.conf or lilo.conf:

```
grub
```

      or, if boot loader is lilo

```
lilo
```

9. Reboot the system using the *New boot* menu option.
10. Verify that sd_mod, sg, mpp_Upper, qla2300, and mpp_Vhba were loaded after the reboot:

```
/sbin/lsmod
```

11. Verify that the RDAC driver discovered the available physical LUNs and created virtual LUNs for them:

```
ls -lR /proc/mpp
```

You can now issue I/Os to the LUNs.

Continue with "Configuring RDAC."

## Configuring RDAC

This task provides instructions for configuring the RDAC multipath devices, which can be used to verify a correct installation of the RDAC software.

You must have completed the installation of RDAC if you are using DS4000 (FAStT) technology.

After you have completed the installation of the RDAC multipathing software: View your RDAC devices by entering the following command:

```
SMdevices
```

# Chapter 9. Uninstalling SAN File System

This topic describes 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, for an AIX 5.2 system:

   ```
   installp -u sfs.client.aix52
   ```

## 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 10. Upgrading 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 version 2.2.

It is assumed that you have applied any required patches that were made to version 2.1 before you upgrade to version 2.2.

Use the following checklist to upgrade from version 2.1. Before you begin upgrading the SAN File System, make sure that you have access to the SAN File System CD-ROM.

Also, see "Obtain prerequisite software" on page 46 for a list of the prerequisite software that is not provided on the SAN File System CD-ROM and which you will need to obtain for upgrading from version 2.1 to 2.2. Once you have obtained the software, follow the relevant installation instructions as described in this guide.

| Steps | | For more information... |
|---|---|---|
| 1 | Create a backup of your existing system. | Chapter 11, "Backing up the SAN File System," on page 125 |
| 2 | Upgrade the master console. | *Master Console User's Guide* |
| 3 | Upgrade the Linux kernel. | "Upgrading the Linux kernel" on page 53 |
| 4 | Upgrade the package repository on each metadata server engine in the cluster. | "Upgrading the package repository" on page 109 |
| 5 | Upgrade the master metadata server engine. | "Upgrading metadata server engines" on page 110 |
| a | Prepare the engine for upgrade. | "Preparing the metadata server for upgrade" on page 110 |
| b | Upgrade the administrative server package. | "Upgrading the administrative server package" on page 112 |
| c | Upgrade the metadata server package. | "Upgrading the metadata server package" on page 112 |
| d | Restart the metadata server engine. | "Restarting the metadata server engine" on page 113 |
| 6 | Upgrade each subordinate metadata server engine. | "Upgrading metadata server engines" on page 110 |
| a | Prepare the engine for upgrade. | "Preparing the metadata server for upgrade" on page 110 |
| b | Upgrade the administrative server package. | "Upgrading the administrative server package" on page 112 |
| c | Upgrade the metadata server package. | "Upgrading the metadata server package" on page 112 |
| d | Restart the metadata server engine. | "Restarting the metadata server engine" on page 113 |
| e | At this point, you can commit the upgrade, or wait and commit the upgrade after you have upgraded the clients as described in the following steps. | "Committing the upgrade" on page 113 |

| Steps | | | | | For more information... |
|---|---|---|---|---|---|
| 7 | Upgrade clients.<br><br>**Note:** Upgrading a client from Version 2.1 to Version 2.2 is optional (that is, a 2.2 metadata server works with both a 2.1 and a 2.2 client). However, if you need Version 2.2 features, such as heterogeneous security, then upgrading clients to Version 2.2 client is required. | | | | |
| | a | Upgrade Windows clients. | | | "Upgrading SAN File System on a Windows client" on page 114 |
| | | 1 | Obtain Windows client software. | | "Obtain version 2.2 software for a Windows client" on page 74 |
| | | 2 | Prepare the client for upgrade. | | "Preparing a Windows client for upgrading" on page 114 |
| | | | a | Stop all client applications. | "Stop all applications on the client" on page 115 |
| | | | b | Uninstall the Windows client | "Uninstalling the SAN File System software from a Windows client" on page 103 |
| | | 3 | Install the software. | | "Installing the SAN File System software on a Windows client" on page 74 |
| | | 4 | Validate client installation. | | "Validating the installation of SAN File System on a Windows client" on page 75 |
| | | 5 | Set up the client to automatically start when the client is rebooted. | | "Automate client restart on reboot" on page 75 |
| | b | Upgrade AIX clients. | | | "Upgrading SAN File System on an AIX client" on page 116 |
| | | 1 | Obtain AIX client software. | | "Obtain version 2.2 software for an AIX client" on page 76 |
| | | 2 | Prepare the client for upgrade. | | "Preparing an AIX client for upgrading" on page 117 |
| | | | a | Stop all client applications. | "Stop all applications on the client" on page 117 |
| | | 3 | Install the software. | | "Installing the SAN File System software on an AIX client" on page 77 |
| | | 4 | Validate client installation. | | "Validating the installation of SAN File System on an AIX client" on page 80 |
| | c | Upgrade Linux clients. | | | "Upgrading SAN File System on a Linux client" on page 118 |
| | | 1 | Obtain Linux client software. | | "Obtain version 2.2 software for a Linux client" on page 80 |
| | | 2 | Prepare the client for upgrade. | | "Preparing a Linux client for upgrading" on page 119 |
| | | | a | Stop all client applications. | "Stop all applications on the client" on page 119 |

| Steps | | | For more information... |
|---|---|---|---|
| | 3 | Install the software. | "Upgrading the SAN File System software on a Linux client" on page 120 |
| | 4 | Validate client installation. | "Validating the installation of SAN File System on a Linux client" on page 82 |
| d | | Upgrade Solaris clients. | "Upgrading SAN File System on a Solaris client" on page 120 |
| | 1 | Obtain Solaris client software. | "Obtain version 2.2 software for a Solaris client" on page 83 |
| | 2 | Prepare the client for upgrade. | "Preparing a Solaris client for upgrading" on page 121 |
| | a | Stop all client applications. | "Stop all applications on the client" on page 121 |
| | 3 | Install the software. | "Installing the SAN File System software on a Solaris client" on page 83 |
| | 4 | Validate client installation. | "Validating the installation of SAN File System on a Solaris client" on page 85 |
| 8 | | Upgrade the cluster. | "Committing the upgrade" on page 113 |
| 9 | | Back up the complete system. | Chapter 11, "Backing up the SAN File System," on page 125 |

## 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, keep all copies of the package repository 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. Insert the SAN File System CD-ROM into the CD-ROM drive, and then mount CD-ROM drive.

   ```
   mount /media/cdrom
   ```

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 name of the currently installed package repository.

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

4. Install the new package repository:[1]

   ```
   /media/cdrom/install/install_sfs-package-2.2.0-104.i386.sh
   ```

5. When prompted for your preferred language, type the number that corresponds to your preferred language and press **Enter**.
6. When prompted to view the International Program License Agreement press **Enter**.
7. After reading and agreeing to the license (by pressing **Enter** to page forward and typing **99** and pressing **Enter** to page backwards), type **1** and press **Enter** to accept the license agreement and install the software.

Continue with the next metadata server in the cluster until you have upgraded the package repository on all metadata server engines.

## Upgrading metadata server engines

This topic lists the steps for upgrading metadata server engines from version 2.1.

Complete the following procedures for all metadata servers in the cluster. Upgrade the master metadata server engine first, and then upgrade the subordinate engines:

1. Determine whether IBM Director Agent is installed, and uninstall it if it is.
   a. Issue the following command:

      ```
      rpm -qa | grep Agent
      ```

      The following results are returned when IBM Director Agent is installed:

      ```
      ITDAgent-4.11-1
      DirAgent-4.11-1
      ```

   b. If IBM Director Agent is installed, uninstall it with the following commands:

      ```
      /etc/init.d/dacmimom stop
      /etc/init.d/TWGagent stop
      rpm -e DirAgent-4.11-1
      rpm -e ITDAgent-4.11-1
      ```

2. Back up the existing server configuration files to a temporary directory. You can also create a checklist similar to the one in "Linux client upgrade checklist" on page 119 to assist you in upgrading the metadata server engines.
3. Prepare the server to be upgraded.
4. Upgrade the administrative server package on the engine.
5. Upgrade the metadata server package on the engine.
6. Verify the upgrade process and restart the engine.

### 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. If the new package repository name is different than the existing package repository, remove the existing package repository.

   ```
   rpm -e existing_package_repository_name
   ```

You can also prevent the metadata server from restarting automatically when it is down by issuing the **stopautorestart** command as described in "stopautorestart."

To maximize the availability of the SAN File System cluster in the event of a power outage, ensure that the following conditions are met before you prepare the metadata server for upgrade:

- A correctly installed SAN File System configuration requires that each power supply within a SAN File System engine be powered by separate and independent power circuits. When the engines are connected in this manner, there can be no single point of failure in the configuration.
- Of the two power supplies in each engine, at least one of these should be powered through an uninterruptible power supply (UPS).
- If you do not use a UPS, then use the RSA II's external power supply and connect it to a UPS as described below.

The external power supply for the RSA II adapter provides access to the RSA II event log if power to the metadata server engine (in which the RSA II adapter is installed) is not available. Perform the following steps to connect the power-supply power cord to your server:

1. Connect the power-supply cable to the external power-supply connector on the RSA II.
2. Connect one end of the power cord to the connector on the power-supply adapter.
3. Connect the other end of the power cord to an uninterruptible power supply.
4. Verify that the RSA II power LED is lit, which means that there is power to the RSA II

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:

   ```
   /usr/tank/admin/bin/sfscli lsserver
   ```

3. Stop the metadata server that you are going to upgrade:

   ```
   /usr/tank/admin/bin/sfscli stopserver server_name
   ```

4. When prompted to confirm that you want to stop this server, respond with `yes`.
5. Verify that the server has stopped:

   ```
   /usr/tank/admin/bin/sfscli lsserver server_name
   ```

   The selected server should be listed as `Not Running`.
6. Optional: If you are using a heterogeneous file sharing feature, obtain and install the Samba Winbind update from www.samba.org.

## stopautorestart

Disables the metadata server from restart automatically if it is down.

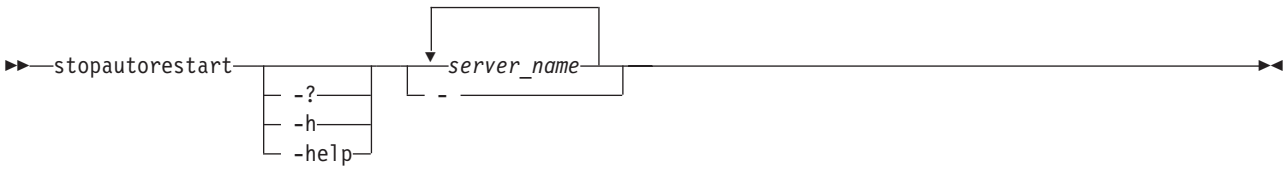

**Parameters**

**–? | –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.

*server_name*
Specifies the names of one or more metadata servers to disable from restarting automatically.

**–** Specifies that, in single-shot command mode, this command is to receive from the input stream (stdin) the names of the metadata servers to disable from restarting automatically.

**Prerequisites**

You must have Administrator privileges to use the command.

**Description**

**Note:** If you run this command from an engine hosting a subordinate metadata server, you can stop the metadata server restart service on only the local metadata server. If you run this command from the engine hosting the master metadata server, you can stop the metadata server restart service on any metadata server.

**Example**

**Disable the automatic-restart service** The following example disables the automatic-restart service for metadata server ST1:

```
sfscli> stopautorestart ST1
The automatic restart service for server ST1 successfully disabled.
```

## 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.

To upgrade the metadata server package, enter the following command:

```
rpm -Uvh metadata_server_package_name
```

For example:

```
rpm -Uvh /usr/tank/packages/sfs-server-linux
```

**Attention:** Running rpm -e instead of rpm -U when upgrading the metadata server might cause problems when restarting, as described in the following bullets:

- If you did run the **rpm -e sfs.admin.linux** command, some configuration files and settings might be lost. Verify that the following entry exists in /etc/inittab before attempting to restart the metadata server:

  ```
  sfs:35:wait:/etc/rc.d/init.d/sanfs start
  ```

- If the entry is missing, run the **setupsfs** command without the **-overwrite** option to restore it. If you use the -overwrite option with the setupsfs command, **all data is lost**.

## Restarting 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:

   ```
   /usr/tank/admin/bin/sfscli startserver upgraded_metadata_server_name
   ```

3. When prompted to confirm that you want to start this engine, respond with `yes`.

## Committing the upgrade

This topic describes how to commit a cluster upgrade.

Before you commit the upgrade, ensure that all metadata servers in the SAN File System cluster have been upgraded. Enter the following command to see the installed version of all servers in the cluster:

```
sfscli lsserver -l
```

The version numbers of all servers in the cluster appear in the last field of the screen output (Most Current Software). Verify that all the version numbers match. If they match, commit the upgrade with the following steps:

1. Commit the upgrade:

   ```
   sfscli upgradecluster
   ```

2. When prompted whether you want to commit the upgrade, enter ***y*** to confirm that you want to commit.
3. Verify that the software version and committed software version values match:

   ```
   sfscli statcluster
   ```

   You see the values for the versions in the screen output. The values for Software Version and Committed Software Version must match.

## 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 2.2 software for a Windows client" on page 74.
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."
3. Install the latest version of the client software. See "Installing the SAN File System software on a Windows client" on page 74.
4. Validate the installation of the client software. See "Validating the installation of SAN File System on a Windows client" on page 75.

### 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 2.2 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.

## Upgrading the SAN File System software on a Windows client

This topic describes how to install version 2.2 of the SAN File System on a Windows client.

- The client for Windows can be installed only on a Windows 2000 server, Windows advanced server, or Windows 2003 standard server. A minimum version 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 is 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**.

# 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 2.2 software for an AIX client" on page 76.
2. Prepare the AIX client for upgrading by stopping all applications on the client. See "Preparing an AIX client for upgrading" on page 117.
3. Install the latest version of the client software. See "Installing the SAN File System software on an AIX client" on page 77.
4. Validate the installation of the client software. See "Validating the installation of SAN File System on an AIX client" on page 80.

## 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 2.2 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 exists 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 -cuxV /mnt/sanfs
   ```

   Example:

   ```
   aix:/mnt/sanfs/sanfs >fuser -cuxV /mnt/sanfs
   /mnt/sanfs:
   vfs    fd=3       21468(root)
   vfs               22742c(root)
   ```

   **Restriction:** 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
   ```

## Upgrading SAN File System on an AIX client

This topic describes how to install version 2.2 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.
- For AIX 5.3, the bos.mp, bos.up, or bos.mp64 must be at least 5.3.0.0 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 see this choice only if you are enabling asynchronous I/O on AIX 5.2 or AIX 5.3. 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. At this point, you can select **Configure Defined Asynchronous I/O** in the smit menu to apply the changes and then exit smit, or you can exit smit and 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 are prompted to enter values for the client configuration, as shown in Table 7 on page 78. In most cases you can accept the defaults. If you change these values, make sure that you type the new values correctly.

## 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. Upgrade the Linux kernel.

2. Obtain the latest version of the SAN File System client software. See "Obtain version 2.2 software for a Linux client" on page 80.
3. Prepare the Linux client for upgrading by stopping all applications on the client. See "Preparing a Linux client for upgrading."
4. Install the latest version of the client software. See "Upgrading the SAN File System software on a Linux client" on page 120.
5. Validate the installation of the client software. See "Validating the installation of SAN File System on a Linux client" on page 82.

## 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 2.2 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 -um /mnt/sanfs
   ```

   Example:

   ```
   lin:/mnt/sanfs/sanfs >fuser -um /mnt/sanfs
   /mnt/sanfs:          14775c(root) 15088c(root) 15089 15089c(root) 15090c(root)
   ```

   **Restriction:** 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 uses the configuration information stored in /usr/tank/client/config/stclient.conf.

## Upgrading the SAN File System software on a Linux client

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_RHEL2.2.1-n
   ```

   or

   ```
   rpm -U sfs.client.linux_SLES82.2.1-n
   ```

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 are prompted to enter values for the client configuration, as shown in Table 8 on page 81.

   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 2.2 software for a Solaris client” on page 83.
2. Prepare the Solaris client for upgrading by stopping all applications on the client. See “Preparing a Solaris client for upgrading” on page 121.
3. Install the latest version of the client software. See “Installing the SAN File System software on a Solaris client” on page 83.
4. Validate the installation of the client software. See “Validating the installation of SAN File System on a Solaris client” on page 85.

## 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 2.2 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 104.
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 -cu /mnt/sanfs
   ```

   Example output:

   ```
   sol:/mnt/sanfs/sanfs >fuser -cu /mnt/sanfs
   /mnt/sanfs:     9823c(root)    9820c(root)    9819o(root)
   ```

   **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.

## Upgrading the SAN File System software on a Solaris client

This topic describes how to install version 2.2 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
   ```

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 are prompted to enter values for the client configuration, as shown in Table 10. In most cases, you can accept the defaults.

*Table 10. Solaris client configuration prompts*

| Parameter | Default | Description |
|---|---|---|
| devices | pat=/dev/dsk/c*t*d*s2 | The SAN File System client determines which disks to use as SAN File System user data volumes by searching a list of disks, called device candidates. The device candidate list can be specified as a pattern or directory. pat=<*pattern*> dir=<*directory path*>. |
| convertertype | ISO-8859-1 | The NLS converter provides the metadata server with data on how to convert strings from the SAN File System client into Unicode. Refer to the International Components for UNICODE Web site noted in "Web sites" on page vii for a list of supported converters. |
| clientname | hostname | The name can be any string, but must be unique among all SAN File System clients. |
| server_ip | No default | The SAN File System client must connect to one of the metadata servers in the cluster. After the client establishes a connection to the server, the server notifies the client of any other servers in the cluster. |
| server_port | 1700 | The SAN File System client must connect to the client-server port on the metadata server. Running the command **sfscli statserver –netconfig <*server_name*>** on the server displays the client-server port. |

*Table 10. Solaris client configuration prompts (continued)*

| Parameter | Default | Description |
|---|---|---|
| mount_point | /mnt/sanfs | The client setup utility mounts the SAN File System to a specified mount point (directory) and creates the file system image. If the specified mount point does not exist, it is created.<br><br>**Attention:** Do not enter the mount point or directory of any general directories that are used by the base operating system of the client. For example, /, /root, /var, /etc, /usr. Doing this might cause the client operating system to stop performing basic functions. If you do mount the client at a standard directory, call the IBM Support Center for assistance. |
| readonly | No | If you mount the SAN File System as read-only, data and metadata in the file system can be viewed, but not modified. |

# Chapter 11. 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 126 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 128 for saving additional metadata server and client configuration files.

**Backing up filesets**

Go to "Saving a FlashCopy image of a fileset and accessing it" on page 129 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 a 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. The metadata server and all clients must complete all active transactions and save their data to disk.
2. From the master metadata server, use the following command to quiesce the SAN File System metadata servers:

   ```
   /usr/tank/admin/bin/sfscli quiescecluster -state full
   ```

   This procedure also locks out any subsequent new input and output from the clients or metadata servers.
3. Start the storage subsystem copy service using the procedure defined in its accompanying documentation.
4. After the storage subsystem copy is complete, use the following command to restart the SAN File System metadata servers:

   ```
   /usr/tank/admin/bin/sfscli resumecluster
   ```
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/sfscli 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 and accessing it" on page 129.

3. Restart the client applications using the specific procedures for those applications.
4. From the master metadata server engine, save the most recent metadata to accompany the FlashCopy images.

   ```
   /usr/tank/admin/sfscli 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).
5. On each metadata server, back up 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). The list of files to be backed up is specified in /usr/tank/admin/config/backup.list. You can customize the contents of backup.list to save any files of your choosing into the DR tarball.
6. Use the backup application to back up the *.flashcopy* directories. Also, back up the file containing the metadata (the file created by the **mkdrfile** command) and the *tar.gz file created by setupsfs -backup.

   **Attention:** Back up Microsoft Windows filesets only from a Windows client; back up AIX filesets only from an AIX client.

- 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.

  However, if your existing backup application does not provide the features required for the enhanced method, follow this 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. From the master metadata server engine, save the most recent metadata.

     ```
     /usr/tank/admin/bin/sfscli 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 are not 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 136.

### 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.

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. 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.

Each system provides unique log information that 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.

1. Use the following steps for metadata servers:
   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
      ```
2. Use the following steps for UNIX-based clients:
   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
      ```
3. Use the following steps for Windows clients:
   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
      ```

**Tip:** Output is stored in the *C:\Documents and Settings\Administrator\Application Data\IBM\Storage Tank\OBDC* directory.

b. Follow the menu options to collect the data. At the C:\WINNT> prompt enter:

```
obdc
```

## Saving a FlashCopy image of a fileset and accessing it

FlashCopy images are saved on a per-fileset basis.

- Using the GUI, select **Manage Copies** in the My Work pane, and then select **Create FlashCopy Images**.
- Click **Next**.
- Under **Select Filesets**, 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, and then access, the FlashCopy images including the newly created image, select **Manage Copies** and then **FlashCopy Images**. The default FlashCopy image name (Image-1) is 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 12. SAN File System installation commands

This topic provides an overview of the SAN File System installation commands.

The SAN File System installation commands include:

- **migratedata**: used to migrate data to SAN File System.
- **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*.

## migratedata

Migrates data to SAN File System.

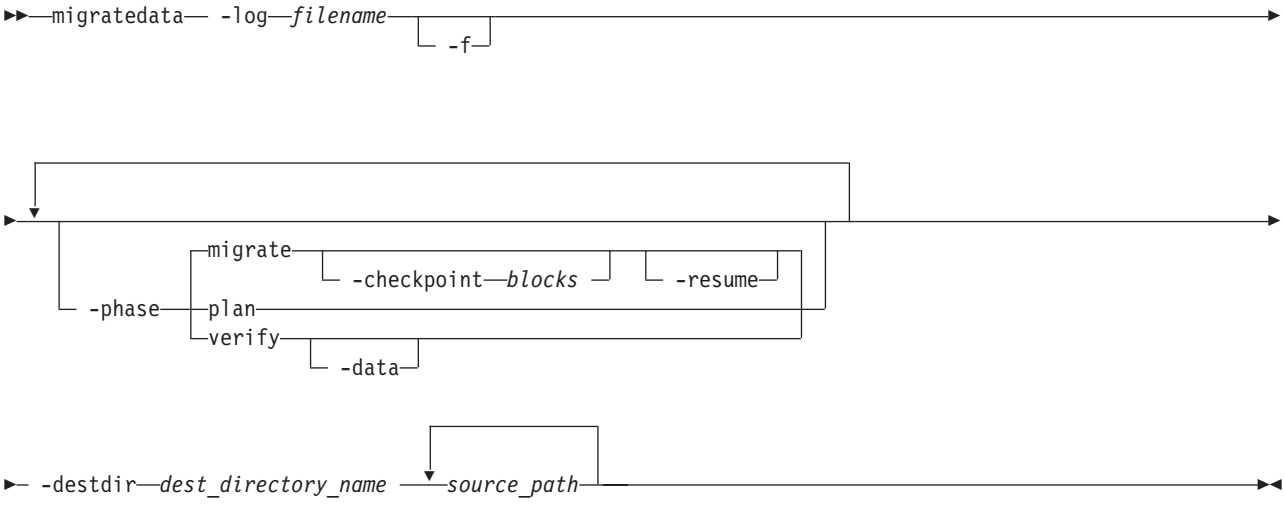


**Parameters**

**–log** *filename*
Specifies the location of a file in which to log migration activities, warnings, and errors. When used with the **–plan migrate –resume** parameter, the **–log** parameter specifies the log file from which to read information about the last completed block or file.

**Attention:** You must specify the correct log file with migrate **–resume** and verify that the source and destination directories specified on the command line match those in the log file.

- If you specify an incorrect log file and the **–f** parameter, **–resume** displays a warning, but migrates all of the data that you specify. However, the verification fails if any of the source directories specified are different than those listed in the incorrect log file.
- If you specify an incorrect log file, but do not specify the **–f** parameter, this command displays an error and exits.

**–f** Specifies that the migration should continue even if there is an error with a file. If specified with the **–phase migrate** parameter, this command skips any files with errors, and continues with the migration process. Examples of file errors include insufficient privileges to read the file, or not running as superuser, preventing the permissions on the SAN File System file to be set. If not specified, an error results in the entire migration being stopped before the file that caused the error. You can then restart the migration after fixing the error.

If specified with the **–phase verify** parameter, this command adjusts any missing metadata attributes, such as permissions and times. If there is a mismatch in size, however, this command does not try to readjust the metadata attributes.

**–phase**
Specifies the migration phase to run. Choices include:

**plan** Gathers information about the available system resources (available memory, number of CPUs, size of the source tree and space available on the destination file system), copies sample files from source directory to estimate transfer rates, and provides an estimated time for the migration of the data set. The copies of the sample files are then deleted, unless the process is interrupted, in which case the copies are not deleted.

**migrate**
Reads data from source file system and writes the data to the destination file system. Although not required, for large data sets, you should run this command in planning mode first. You can stop the migration process at any point (by pressing **Ctrl+C**) and resume from the last completed file or block (using the **–resume** parameter).

This is the default value.

**verify** Verifies the integrity of the migrated data using the Message Digest 5 verification algorithm on the contents of the file, as well as verifying consistency of the metadata (such as owner and modification time stamp settings) between the source and destination files.

You can specify more than one phase. For example, to plan, migrate, and verify the data, specify **–phase plan –phase migrate –phase verify**. Although you can specify the phases in any order, this command always estimates the completion time, migrates data, and then verifies the migrated data.

If the **–phase** parameter is not specified, this command runs only the migration phase.

**–checkpoint** *blocks*
Shows the progress when migrating large files. If you specify this parameter, the **migratedata** command writes a checkpoint in the log file after each

specified number of blocks of a file has been migrated. (The block size depends on the client platform.) For example, if you specify **–checkpoint 20**, this command makes an entry in the log file each time 20 blocks of file data is migrated. On a platform with a block size of 16 MB, this command writes to the log file after each 320 MB of data from a file has been migrated.

If the migration process is interrupted, this parameter allows you to resume the migration at the place it left off.

If unspecified, the **migratedata** command makes an entry in the log file after each complete file has been migrated. You can resume the migration at the point of the last migrated file.

**–resume**
Resumes the migration from the last completed block or file (logged in the log file specified by the **–log** parameter). If the log file indicates that some files in the source directory are migrated and this parameter is not specified, this command restarts the migration process from the beginning (performs a fresh migration).

**–data**
Verifies every block of source data (file data and metadata) with the destination data. If not specified, this command verifies only the metadata unless there is a mismatch in the file attributes, in which case this command then verifies the file data.

**Note:** Verifying all data is very time consuming and can take as long as the migration itself.

**–destdir** *dest_directory_name*
Specifies the name of the destination directory for the migrated data. The directory can either exist or be a new directory name. It is recommended that you create the directory before beginning the migration process. If the directory does not exist, this command creates the directory using the default permissions.

*source_path*
Specifies one or more paths of directories or files to migrate.

**Prerequisites**

You must have root privileges on a UNIX-based client or Administrative privileges on Windows to use this command.

All storage pools, all filesets, and at least one policy must be set up. All activity (from applications, such as database servers and application servers, or users) that modifies data on the source and destination file systems must be stopped and remain stopped to guarantee consistency of the migrated data.

The destination directory must exist with correct set of permissions and appropriate storage policies must be configured.

**Example**

**Migrating data** This example migrates data from the work/capital directory on the client machine to the sanfs/cnt1 directory in the global namespace. A checkpoint is written to the mgrt_capital.log log file each time 20 blocks of file data is migrated.

```
migratedata -log /mgrtlogs/mgrt_capital.log -phase migrate -checkpoint 20
-destdir /mnt/tank/sanfs/cnt1 work/capital
```

## setupsfs

Configures and starts a SAN File System metadata server.

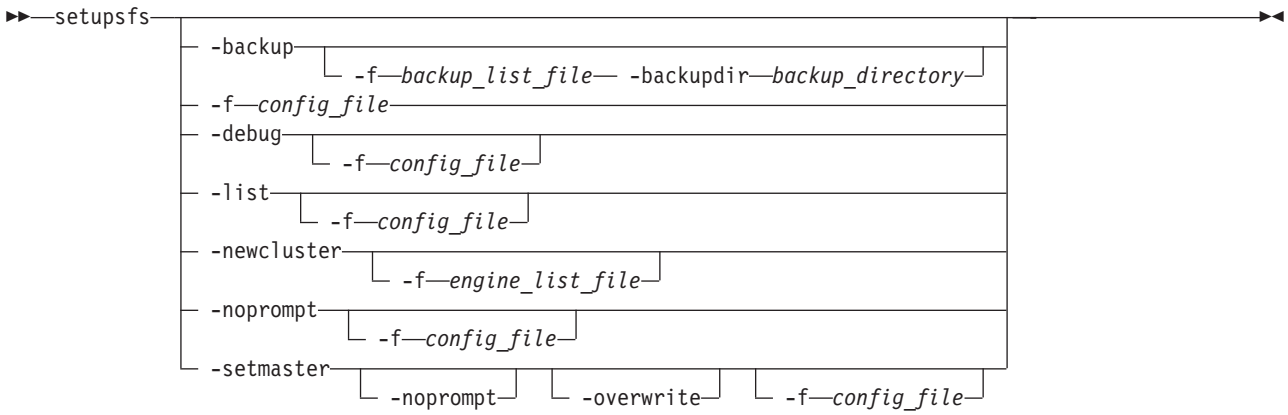


**Parameters**

**–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. You can also customize the contents of backup.list to save any files of your choosing into the DR tarball.

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 -xzvf /usr/tank/server/DR/DRfiles-hostname-date.tar.gz
```

Then, run **setupsfs** without parameters (or with only the **–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 Active Directory 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 destroys 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 are presented as the suggested defaults when the prompt is displayed. Otherwise, the manufacturing defaults are presented.

## setupstclient

Configures and starts SAN File System clients.

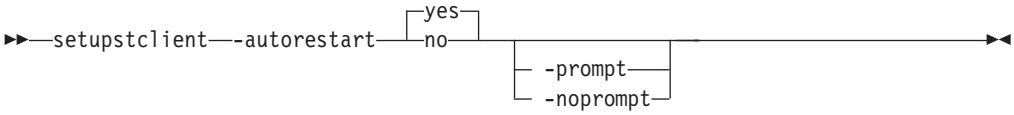


**Parameters**

**–autorestart**
Specifies whether to restart the SAN File System client automatically at boot time.

**yes** Enables autorestart. The SAN File System client is restarted automatically at boot time. This is the default value.

**no** Disables autorestart.

**–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:

```
/usr/tank/client/bin/setupstclient
```

## tmvt

Validates that the existing hardware and software on a metadata server meet all installation requirements for the SAN File System.

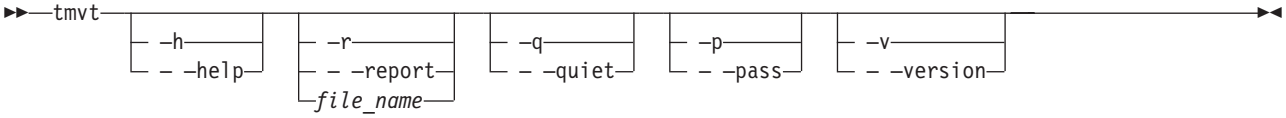


**Parameters**

**–? | –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 quite 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

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.

**Note:** If the kernel level on which you are running is not supported by SAN File System, you should correct the kernel version and rerun tmvt.

### 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            1.1.7            1.1.7
              Machine BIOS Build               GEJ57B           GEJ57B
                 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 can find a wide variety of sources available from IBM to assist you.

Services available and telephone numbers listed are subject to change without notice.

**Software Maintenance Agreement**

All distributed software licenses include Software Maintenance Agreement(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. You can elect to extend coverage for a total of three years from date of acquisition. While your Software Maintenance is in effect, IBM provides you assistance for your 1) routine, short duration installation and usage (how-to) questions; and 2) code-related questions. IBM provides assistance by 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

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 11 lists some of these pages.

*Table 11. 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 |

## 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 can 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.

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.[2]

2. Response time varies 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. Troubleshooting the installation

This topic describes technical problems that might occur during the installation of SAN File System.

The following topics describe the problems and provide solutions that help you troubleshoot the installation process.

## Finding and correcting problems

This topic describes items to check to determine and correct problems in SAN File System.

1. In the tank.properties files, the **RSA_USER** and **RSA_PASSWD** entries must be the same on all servers. If the values do not match, correct the values using **vi editor** and save the file. Be careful not to space after the last letter.
2. Run the following commands: **stopCimom**, **startCimom**, **stopConsole**, **startConsole**. These commands are located in /usr/tank/admin/bin.
3. Retry the **tanktool lsengine** command to see if connectivity is restored. The default values are: **RSA_USER=USERID** and **RSA_PASSWD=PASSW0RD** (zero not the letter O). You can change the values, but make the changes on all servers.
4. If the **RSA_USER** and **RSA_PASSWD** information match on all servers, try to log into each RSA card user interface using the userid and password in the tank.properties files. From the Configuration File of each RSA card, you can view the Login Profile entries to see if there is an entry to match the data in the tank.properties files.
5. If you cannot log in to a server or find a matching entry, add a new login profile.
6. As a test, log out of the RSA user interface and back in using the new Login information.
7. Run the following commands: **stopCimom**, **startCimom**, **stopConsole**, **startConsole**. These commands are located in /usr/tank/admin/bin.
8. Retry the **tanktool lsengine** command to see if connectivity is restored.

   **Note:** Multiple Login profiles can exist, but at least one login profile must be the same on all servers; and each must match what is in all the tank.properties files. The Host operating system entry in the Configuration Summary File must say Linux. If it does not, the `In Boot`, `Booting OS`, and/or `Wrong OS` errors might appear and the *ibmasm* daemon might not start.
9. To correct, in the left panel, select Server->ASM Control->System Settings->ASM Information.
10. Select **Linux** from the Host operating system drop-down menu.
11. In the left panel, select **Restart ASM**.
12. Run the following commands: **stopCimom**, **startCimom**, **stopConsole**, **startConsole**. These commands are located in /usr/tank/admin/bin.
13. Retry the **tanktool lsengine** command to see if connectivity is restored.
14. If the command still gives an error, reboot the metadata server to start/restart the ASM daemon on the RSA card.

15. Verify that the RSA cabling and configuration is correct. You can verify the information using the RSA Configuration Summary File. Refer to the *SAN File System Planning, Installation, and Configuration Guide for more information*.

## Supplying power to metadata server engines

This topic describes information about supplying power to metadata server engines.

**Problem:** Because the external Remote Supervisory Adapter II (RSA) power is not required to ensure SAN File System metadata server redundancy, you must ensure that the power supply units for each metadata servers' engine are connected to separate and independent power supply circuits.

**Solution:**To properly install and configure SAN File System, each power supply within a metadata server's engine must have a separate and independent power circuit. If servers are connected in this manner, there cannot be a single point of failure in the configuration and no external RSA II power is necessary. The external RSA II power supply is not a SAN File System requirement for high availability. If you want additional power-supply redundancy, use a third independent circuit for the RSA II. You can order an RSA II power supply in addition to SAN File System to connect the RSA II and the third independent circuit. For additional instructions on installing SAN File System metadata server engines, refer to the *SAN File System Planning, Installation, and Configuration Guide*.

# Appendix F. Troubleshooting the RSA II adapter

This topic provides information that you can use to resolve problems with the RSA (Remote Supervisor Adapter) II adapter.

Use the following topics to determine and resolve problems with the RSA II adapter.

## Configuring an IP address for each RSA II

This topic provides information about configuring an IP address for each RSA II.

**Problem:** Remote Supervisory Adapter IIs (RSA) are installed, but not configured in each SAN File System metadata server engine used in a SAN File System cluster.

**Solution:** Configure each RSA II with its own IP address, independent of the host metadata server and other RSA IIs. Ensuring that each RSA II is properly configured enables better monitoring and control of each metadata server in the SAN File System cluster. Proper configuration also provides redundant communication paths within the cluster.

For complete instructions for properly configuring the RSA II, refer to the *SAN File System Planning, Installation, and Configuration Guide*.

## RSA II adapter errors during installation or upgrade

This topic describes Remote Supervisory Adapter (RSA) II-related errors with X-series hardware during a SAN File System installation or upgrade.

**Problem:** During a clean install of the SAN File System or during an upgrade to SAN File System from a previous release, you may encounter errors related to the RSA II which is part of the IBM X-series platform. Some examples of these errors include:

- `CMMNW5072E` accessing an engine in the user interface
- `HST550030` or `HST550031` during **setupsfs** execution
- various `ibmasm` errors

**Cause:** Possible communication failure with the RSA II.

**Solution:** Examine the TMVT (Target Machine Verification Tool) reports for each metadata server. If you find entries in the *Failed Hardware Component Checks* for the RSA II that show ″missing,″ the RSA II might be failing. Try reseating the RSA II on the metadata server with the missing entry, and then rerunning the TMVT report. If the RSA II still shows ″missing,″ contact IBM hardware support.

# Appendix G. 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 106-0032, 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 | Tivoli | TotalStorage |
| WebSphere | xSeries | |

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.

Linux is a trademark of Linus Torvalds 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. 

# Index

## Numerics



## A



## B



## C



## D



## E



## F



## G



## H



## I



## J

## V



## W



## X



## Z

# 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 2**

**Publication No. GA27-4316-02**

**Overall, how satisfied are you with the information in this book?**

| | Very Satisfied | Satisfied | Neutral | Dissatisfied | Very Dissatisfied |
|---|---|---|---|---|---|
| Overall satisfaction | ☐ | ☐ | ☐ | ☐ | ☐ |

**How satisfied are you that the information in this book is:**

| | Very Satisfied | Satisfied | Neutral | Dissatisfied | Very Dissatisfied |
|---|---|---|---|---|---|
| Accurate | ☐ | ☐ | ☐ | ☐ | ☐ |
| Complete | ☐ | ☐ | ☐ | ☐ | ☐ |
| Easy to find | ☐ | ☐ | ☐ | ☐ | ☐ |
| Easy to understand | ☐ | ☐ | ☐ | ☐ | ☐ |
| Well organized | ☐ | ☐ | ☐ | ☐ | ☐ |
| Applicable to your tasks | ☐ | ☐ | ☐ | ☐ | ☐ |

**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. _____

**Readers' Comments — We'd Like to Hear from You**
GA27-4316-02

Cut or Fold
Along Line

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

GA27-4316-02

Cut or Fold
Along Line

Printed in USA

GA27-4316-02