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

Planning, Installation and Configuration Guide

Version 1 Release 1

GA27-4316-00

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

Planning, Installation and Configuration Guide

Version 1 Release 1

GA27-4316-00

NOTE

Before using this information and the product it supports, read the safety information in "Safety information" on page 125 and the general information in Appendix J, "Notices", on page 123.

First Edition (November 2003)

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 guide provides information useful to planning, installing and configuring IBM TotalStorage SAN File System. The information is organized as follows:

- Chapter 1, “Planning”, on page 1 provides you with details about prerequisites and preparation for installing your SAN File System.
- Chapter 2, “Preparing the SAN”, on page 29 describes the steps that you need to perform to prepare your SAN for SAN File System installation.
- Chapter 3, “Physical Installation”, on page 35 provides information about how to perform the physical installation and cabling of SAN File System hardware.
- Chapter 4, “Master console”, on page 43 covers the command line interface, Tivoli SAN Manager, Call Home, and remote access.
- Chapter 5, “Software setup”, on page 51 describes the required steps to configure SAN File System software, including assigning IP addresses, users, clients, setting up zones, and so forth.
- Chapter 6, “Upgrading”, on page 73 tells you how to perform upgrades on the software components of SAN File System.
- Chapter 7, “Backing up”, on page 79 gives step-by-step instructions for performing SAN File System backups in preparation for disaster recovery, using either the LUN or API method. This chapter also explains how to create FlashCopy images for normal backups during full system availability; and how to use the one-button data collection capability.
- Chapter 8, “Uninstalling”, on page 85 describes the procedures for uninstalling SAN File System components.
- The appendices provide the following additional information:
  - Accessibility features of the SAN File System console and help system
  - Getting help
  - Purchasing additional services
  - Basic configuration for a quick start
  - Worksheets related to planning, installating and configuring your SAN File System
  - Disaster recovery information
  - Sample policy sets
  - Managing local drives
  - IBM statement of limited warranty
  - Notices

Who should use this guide

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

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

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

- Knowledge and training in the technology of SAN File System and its functions
- Familiarity with the SAN File System hardware
- Awareness of the procedures in this document
- Awareness of the safety practices in this and related documents
- Awareness of related installation and service publications

Related topics:

Notices in this 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

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* that came with your storage engine and at www.ibm.com/storage/support. To use the CD, insert it in the CD-ROM drive. If the CD does not launch automatically, follow the instructions on the CD label.

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

- *IBM TotalStorage SAN File System Release Notes*

  This document provides any changes that were not available at the time the publications were produced. This document is available only from the technical support Web site: www.ibm.com/storage/support
- *IBM Safety Information — Read This First*, SD21-0030

  This document provides translated versions of general safety notices and should be read before using this product. This document is provided only in hardcopy.
- *IBM Statement of Limited Warranty*

  This publication describes the IBM statement of limited warranty as it applies to the SAN File System Model 1RX storage engine.
- *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 and client command-line interfaces.

- *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 personnel.

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

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

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

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

- *Rack Installation Instructions*

  This publication provides instructions for installing the Model 1RX in a rack.

- *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.

- *Subsystem Device Driver User's Guide for the IBM TotalStorage Enterprise Storage Server and the IBM TotalStorage SAN Volume Controller*, SC26-7540

  The Subsystem Device Driver (SDD) provides the multipath configuration environment support for a host system that is attached to an IBM TotalStorage Enterprise Storage Server® (ESS), IBM TotalStorage SAN Volume Controller or IBM TotalStorage SAN File System. This book provides step-by-step procedures on how to install, configure, and use SDD for the host systems.

  Note: SAN File System supports the version of the Subsystem Device Driver that is shipped with the program product.

- *IBM TotalStorage Translated Safety Notices*, GA67-0043

  This publication contains translated versions of hardware caution and danger statements that appear in the publications in this library. Each caution and danger statement has an assigned number that you can use to locate the corresponding statement in your native language.

Web sites

The following Web sites have additional and up-to-date information about SAN File System:

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

Chapter 1. Planning

This section provides you with details about prerequisites and preparation for installing your SAN File System.

Prerequisites

There are several hardware and software prerequisites that must be confirmed before installing SAN File System.

Prerequisites:

Before installing SAN File System, there are hardware and software prerequisites that must be provided. Lists of requirements and considerations follow:

- Requirements
  - No software may be installed on the metadata servers or the master console that is not part of the preinstalled SAN File System software except for antivirus packages and Virtual Network Computing (VNC). (There is no antivirus requirement for Metadata servers.
  - The SAN configuration should not have a single point of failure. This means that connectivity should be planned to ensure that the loss of an HBA, switch, GBIC, fibre cable, or storage controller can not cause complete loss of connectivity of SAN File System with the storage subsystem.
  - SAN File System metadata servers must have access to all data and metadata LUNs.
  - Separate fabrics should be used for metadata server connectivity. By creating two distinct, non-overlapping paths from each metadata server to storage, more LUNs are possible. If the paths are connected using the same fabric, intra-fabric path overlap would cause four or more paths to be seen by the metadata server, cutting in half the number of useable LUNs.
  - The SAN File System clients must have access to all data LUNs and must be prevented from having access to metadata LUNs.
  - Data LUNs must be able to be shared among hosts that are based on different operating systems.
  - SAN File System metadata server traffic between the SAN network and the storage array must be segregated from the SAN File System client traffic traveling that same path.
  - HBAs for both clients and servers must be isolated from each other to avoid problems associated with them logging in as both hosts and targets.
  - SAN File System fabrics should be isolated from non-SAN File System fabrics on which administrative activities could occur. Administrative activities on switches can affect all users of a switch, even if proper hard zones are present.
- Considerations
  - Client software requirements
  - Number of clients
  - Number of engines per file system
  - Number of file systems per client
  - Size and number of LUNs

  - LDAP server requirements
  - Network requirements, type and connectivity
  - SAN attachment requirements
  - NFS requirements
  - Supported disk subsystems
  - Supported client hardware
  - Application requirements
  - Maximum file size
  - Maximum number of files
- Additional recommendations
  - Client-cluster and intra-cluster comunication traffic are expected to be on the customer LAN.
  - All nodes must be on the same physical network.
  - All hardware engines are expected to be on the same physical network.
  - If multiple subnets are configured on the physical network, it is recommended that all nodes be on the same subnet.
  - The Remote Supervisory Adapter (RSA II) card network port is expected to be on the same physical network.
  - It is recommended that routers and gateways not exist between the clients and the SAN File System server.

Related topics:

- "Data-migration prerequisites" on page 6
- "Hardware prerequisites" on page 3
- "Software prerequisites" on page 4
- Appendix E, "Worksheets", on page 101

Supported browsers

Web browser support:

Most SAN File System functions and online documentation are accessed through standard Web browsers. SAN File System supports the following Web browsers (others may work, but have not been tested):

- Internet Explorer 6.0 and above

  Note: For Internet Explorer 6.0, Service Pack 1 is also needed.
- Netscape 6.2 and above

  Note: While Netscape 6.2 is supported, Netscape 7.0 and above are preferred.

Limitations:

The **Back**, **Forward**, **Refresh** or **Reload** functions of either browser are not supported and may cause unexpected rendering problems. Additionally, opening a hyperlink into a separate browser window is not supported.

Hardware specifications

Dimensions:

Height	85.4 mm (3.36 in.)
Depth	698 mm (27.48 in.)
Width	443.6 mm (17.5 in.)
Weight	21.09 kg to 28.12 kg (46.5 lb to 62 lb)

Operating environment:

Air temperature (at maximum altitude of 2134 m [7000 ft])	10° to 35°C (50° to 95°F)
Humidity	8% to 80%

Nonoperating (power-off) environment:

Air temperature (at maximum altitude of 2134 m [7000 ft])	–40° to 60°C (–40°0 to 140°F)
Humidity	8% to 80%

Heat output:

Minimum configuration	341 Btu/hour (100 watt)
Maximum configuration	2200 Btu/hour (645.2 watt)

Power requirements:

Voltage low range	90 V ac to 137 V ac
Voltage high range	180 V ac to 265 V ac
Power consumption	0.1 kVA to 0.62 kVa
Sine-wave input	50 to 60 Hz required

Acoustical noise emission:

Declared sound power	6.7 bel
Bystander sound pressure	50 dBa

Hardware prerequisites

Verify that the following infrastructure is available for a SAN File System system:

- Adequate rack space is available. Each SAN File System server engine requires 2U of rack space.
- A KVM (Keyboard Video Mouse) available for the server engines. This function is provided when a master console is ordered as part of the system. Otherwise, a KVM must be supplied for each SAN File System server. Optionally, one KDM with an A-B-type switch can be used to allow connection to multiple Linux/Intel-based servers. A third switch is needed for three engines (or one per engine).

  Note:

- A SAN system with two LC switch ports per server engine. The SAN ports should be 2-GB connector running at 1 GB or greater.
- A LAN with a minimum of seven 10/100 Ethernet ports and addresses, (two for each server engine, one each for the RSA II card, and one for the master console).
- IBM TotalStorage Enterprise Storage Server with LUNs defined.
- SAN and LAN cables and GBICs.
- Power outlets (two per server engine).
- SAN-attached clients with supported client operating systems.
- Clients must have IBM-supported fibre channel adapters

Additional requirements:

- Ensure the room air temperature is below 35° C (95° F).
- Do not block any air vents. Usually 15 cm (6 in.) of space provides proper airflow.
- Do not leave open spaces above or below an installed server in your rack cabinet.
- Always install a blank filler panel to cover open space and to help ensure proper air circulation.
- Install your server only in a rack cabinet that has perforated doors.
- Install the heaviest device in the bottom of the rack.
- Do not extend more than one device out of the rack cabinet at a time.
- Remove the rack doors and side panels to provide easier access during installation.
- Connect the server to a properly grounded outlet.
- Do not overload the power outlet when installing multiple devices in the rack cabinet.
- Install your server in a rack that meets the following requirements:
  - Minimum depth between the front mounting flange and the inside of the front door: 70 mm. (2.76 in.)
  - Minimum depth between the rear mounting flange and the inside of the rear door: 157 mm. (6.18 in.)
  - Depth between the front and rear mounting flanges to support the use of the cable management arm: 718 mm. (28.27 in.) to 762 mm. (30 in.)

Related topics:

- "Software prerequisites"
- "Data-migration prerequisites" on page 6

Software prerequisites

The following software is required.

- An SSH client is shipped on the master console and is called PuTTY telnet. SSH connections to the SAN File System engines are required to access to the SAN File System administration CLI.
- LDAP Server. This is required by the Administrative server. See "Gather existing LDAP implementation information" on page 10 for the list of supported LDAP servers.

Note: It is recommended that antivirus software be installed on the master console. This software should be provided and installed by the customer. IBM installers should be certain that the customer is aware of this recommendation. There is no antivirus requirement for Metadata servers.

Prerequisites table:

The following prerequisites must be checked during a given component or platform type-s installation process.

Table 1. Component software prerequisites

Prerequisite	Value	Automated
Metadata server		
Linux Kernel version	2.4 with RPM patches	Yes
Available space	20 MB	Yes
Administrative server		
Linux Kernel version	2.4 with RPM patches	Yes
Available space	20 MB	Yes
SSH support installed		Manual
LDAP server available	See "Gather existing LDAP implementation information" on page 10 for LDAP requirements.	Manual
AIX client		
Operating system version	AIX 5L v 5.1 with maintenance level 4 or greater, 32-bit only	Yes
Available space	20 MB	Yes
SSH installed. This is optional, but highly recommended for clients initiating remote administration.	Any	Manual
FC-HBA drivers installed		Manual
Multipath disk manager. This is optional.	SDD v1.4.0.5	Manual
Windows client		
Operating system version	Windows 2000 Server or Advanced Server, with SP3 or later	Yes
Available space	20 MB	Yes
Networking is installed with IPSEC enabled.	Bundled with Windows.	Manual
FC-HBA drivers installed		Manual
Multipath disk manager. This is optional.	SDD v1.4.0.2	Manual

Related topics:

- "Hardware prerequisites" on page 3

- "Data-migration prerequisites"

Data-migration prerequisites

Verify the following conditions before starting the data-migration utility:

- SAN File System and clients must be installed and properly configured.
- SAN File System must be set up with the appropriate containers, pools, policies, and security.
- The clients must be able to access the source file systems (for example, directly-attached, network-attached storage (NAS), or storage area network (SAN) disks) and SAN File System during the data-migration process.
- When migrating data from a Windows client, you must create the destination directory in the SAN File System file system and verify that the security attributes of the destination directory match that of the source directory. Otherwise, the verification phase of data migration will fail and the migrated data will have incorrect permissions.
- All applications that modify the data being migrated (including database and application servers) must be stopped until the migration completes to guarantee data integrity.
- Twice the space of the data is available for migration. Note that the data-migration utility does not verify whether there is enough space in the system pools where data is being migrated.
- Compressed files are expanded during data migration. Sufficient space must be available in the SAN File System to store the expanded files. Refer to the documentation for your operating system to determine the compression ratio and estimate the amount of space required.
- Sparse files become dense, or full, files during data migration. Sufficient space must be available in the SAN File System to store the dense files.
- To invoke the **migratedata** command, you must either supply the full path or update the PATH environment variable to include the migration directory.

  For clients based on UNIX®, this is done in your shell profile (for example, export PATH=$PATH:/usr/tank/migration/bin).

  For clients for Windows, edit the PATH environment variable (for example, from the Control Panel, double-click **System**, and then click the **Advanced** tab and **Environment Variables**. In the **User variables** group, click the PATH variable and then click **Edit**. At the end of the text in the **Variable Value** field, type: **c:\Program Files\IBM\Storage Tank\Migration**).
- You must have superuser privileges (for UNIX-based clients) or administrator privileges (for Windows clients) to migrate data.

Related topics:

- "Migrating data" on page 71
- "Perform capacity planning" on page 11

Perform site audit

Context: Before SAN File System can be successfully integrated into an existing SAN, you must determine the existing hardware, software, network, and security components.

Steps:

Perform the following steps to perform a site audit:

1. Verify that the system room meets the physical installation requirements.
2. Ensure that adequate space, power connections, and administrative facilities are available to support the physical installation and connectivity requirements of the new or relocated hardware. The following current and voltage ratings are required to power each SAN File System engine:
   - 5 amps @ 100–127 VAC @ 50/60 Hz Single Phase
   - 2.5 amps @ 200–240 VAC @ 50/60 Hz Single Phase

   Each SAN File System engine requires two power inputs, one for each power supply. SAN File System can be installed in any industry standard 19-inch server rack enclosure, such as the IBM 9306 PC Server Rack Enclosure.
3. Gather information about SAN components such as HBAs and switches including vendor, model, firmware revision, and any platform driver software in use. Gather information about volume management software, clustering packages, and multi-pathing software.
4. Determine the zoning configuration parameters. See Chapter 2, “Preparing the SAN”, on page 29.
5. Determine the LUN masking details for access control. See “Determine SAN attachments” on page 12 for more information.
6. Determine the SAN security parameters. (You need security parameters if separate file systems are planned to restrict access to some storage.) See “Determining a security model” on page 20 for more information.
7. Determine the switched zones, if any. This should include:
   - Determining how the switches (fabrics) are zoned
   - Determining if there are separate zones created so that certain hosts cannot see the storage used by another host
8. Determine the LUN masking details, if any, in use on the storage subsystem and what purpose they are used for. See “Determining SAN File System LUN configuration” on page 16 and “Determining which LUNs can be used as SAN File System volumes” on page 17.
9. Determine the cluster software details, including a list of applications and their version levels.

   Note: No other software can use or manage SAN File System storage.
10. Determine the SAN manager details, including the application name and version number.

    Note: IBM Tivoli® Storage Area Network Manager must be ordered with SAN File System. Determine if the functionality of the current SAN manager can be duplicated for the IBM Tivoli SAN Manager.
11. Determine volume manager details, including the application and version level.

    Note: Volume managers, such as Veritas Volume Manager or LVM on AIX, can be used only to manage virtual disks or LUNs that are not managed by SAN File System.

Related topics:

- “Gathering information about SAN components” on page 8
- “Gather information about the existing IP network” on page 9
- “Verify the Fibre Channel SAN” on page 9

Gathering information about SAN components

Steps:

Ensure that the existing SAN components meet the requirements for SAN File System. Specify any components that must be upgraded prior to SAN File System installation. Components to be checked include:

- Client platform operating systems
- IBM fibre channel host bus adapters (HBAs)
- Fibre channel switches or directors
- Storage subsystems firmware
- Drivers
- Hardware

When checking these components, determine the following:

- Operating system versions on each platform
- HBA drivers installed on each platform
- HBA models installed in each platform
- HBA BIOS level
- firmware versions currently installed on each HBA
- Multipathing software being used on each platform
- User applications (especially clustering or volume management software)
- Switch models and their firmware revision levels
- SAN management applications

Note: It is recommended that a drawing of the SAN be provided before the installation begins.

1. Ensure that there are enough switch ports, GBICs, and fibre cables for the SAN File System servers. Each engine has a dual port HBA for SAN access, and each HBA port requires an N_port, on a switch or fabric. Therefore, for a base SAN File System server configuration of two engines, four switch ports are required. Each additional engine requires two more ports. Beyond the switch ports, additional GBICs and fiber cables might be required.

   Note: The SAN File System server HBAs have female LC type fibre cable connectors.
2. Ensure that the disk storage subsystem is supported by SAN File System. Currently, the disk storage subsystems supported for use with SAN File System are:
   - the IBM TotalStorage Enterprise Storage Server, models 2105-F20 and 2105-800
   - the IBM TotalStorage SAN Volume Controller, model 2145
   - the IBM TotalStorage SAN Integration Server, model 2146

   Check the SAN File System Release Notes for the supported code levels of these storage subsystems.
3. Ensure that the SAN has enough ports for the storage subsystems. There should be a minimum of four ports for each storage subsystem for multipathing, failover, and availability, two for metadata server connectivity and at least two for client system connectivity.

Unless the storage subsystems that will be managed by SAN File System are already on the existing SAN, additional SAN ports will be required for this. Documentation or the vendor for the subsystem should be consulted for the number of ports and other SAN planning information. Match SAN component models and versions against the list of components and versions which are supported by the SAN File System. Upgrade versions as necessary.

4. Ensure that the Windows client applications will work with the SAN File System IFS.

Related topics:

- "Verify the Fibre Channel SAN"

Gather information about the existing IP network

Steps:

Perform the following steps to gather information about the existing IP network:

1. Assess the requirements for the interaction of existing customer LANs and the SAN File System client/Metadata server LAN.
2. Determine which hardware modifications, including NICs, cabling, and switches, are needed to support the required LAN topology.

   Note: Your requirements should include administrative LAN connections to new SAN hardware, storage devices, and SAN File System servers.
3. Determine if the existing LAN has enough switch ports for the Metadata servers.

   Each engine has three 10/100 Mb or Gigabit Ethernet interfaces to tie into an existing IP network. There is one IP address for the Metadata server and one for the Remote Supervisory Adapter (RSA) card in each engine.

   The Metadata server configuration includes from two to eight engines. Each engine will have one IP address and one active interface. There are two types of interface: MDS is 10/100/1000 Copper Ethernet, or 1Gb Fibre Ethernet, and RSA is 10/100/1000 Copper Ethernet on the RSA card.
4. Determine which software modifications will be needed to support the resulting LAN topology (clients, servers, and switches).

Related topics:

- "Gathering information about SAN components" on page 8
- "Verify the Fibre Channel SAN"
- "Gather existing LDAP implementation information" on page 10

Verify the Fibre Channel SAN

Context:

Although the SAN can be comprised of heterogeneous switches, each SAN File System fabric (one or more interconnected switches) should be homogeneous. Heterogeneous fabrics have two or more interconnected switches from different vendors, while homogeneous fabrics have one or more interconnected switches from the same vendor. Note that not all HBA-Switch vendor combinations are supported. Check the SAN File System HBA and SAN component compatibility matrix referenced in the Release Note.

Related topics:

- "Perform site audit" on page 6
- "Gathering information about SAN components" on page 8
- "Gather information about the existing IP network" on page 9

Note: Fibre channel Ethernet cables do not ship with the SAN File System hardware.

Gather existing LDAP implementation information

In this section you gather the needed information for implementing LDAP.

Context:

LDAP is used for Administrative servers. You may perform UNIX user mapping, such as authentication and coordination using LDAP.

You must set up and maintain your user database in LDAP using existing LDAP tools. You are also responsible for ensuring that your LDAP is highly available (either by using the existing LDAP server, or setting up the IBM Director Server with replication).

IBM has tested the following LDAP servers:

- IBM Director Server v5.1
- OpenLDAP v2.0 (Linux)

SAN File System administration requires an LDAP server that can be configured for SAN File System. The requirements include:

- You must be able to create four objects under one parent distinguished name (DN), one for each SAN File System role.
- Each role object must contain an attribute that supports multiple DNs.
- You must be able to create an object for each SAN File System administrative user.
- Each administrative user object must contain an attribute that can be used to log in to the SAN File System console or CLI, and a userPassword attribute.

Notes:

1. Installing the LDAP server on a cluster server or on the master console is not recommended and not supported. The LDAP function should be performed by another computer.
2. If LDAP over SSL is used, an authorization certificate must be provided by the customer.

Steps:

Perform the following steps:

1. Get the name of the machine on which the LDAP server is running (and port if not running in the default port normally used by the LDAP server).
2. Decide what schema will be used for storing the information of the SAN file system administrative users and their associated roles.
3. Get an LDAP userID and password that has priviledges to create the entries needed to create users and roles in the LDAP server.
4. Decide what attribute in the roles schema will hold the role name (such as Administrator, and Backup).

5. Decide what attribute in the roles schema will hold the role members or users that are authorized to execute commands under a given administrative role. The sample configuration in this book assumes accessRole.members.
6. Decide which attribute will hold the userID for administrative users.
7. If using a secure LDAP channel, you also need the public certificate (key) from the LDAP administrator so it can be put into the truststore as part of the installation setup procedure

Perform capacity planning

You must determine whether there is enough storage capacity to perform the migration to SAN File System. You also must determine whether there is sufficient storage capacity to provide for some short-term growth in the storage capacity.

Context:

To determine the storage capacity requirements consider the following:

1. The amount of storage space required in the destination SAN File System is at least double the space currently consumed in the source file system. For example, if the data in the source file system occupies *n* blocks, during migration the SAN File System will allocate at least *n* blocks. Therefore, the total storage space required will be more than two times *n* blocks.
2. Estimate about 8 hours to migrate one terabyte of data (which includes about three to four hours for hardware configuration and data verification. The migration of data itself should be about four to five hours for each terabyte of data.
3. Ten percent of the total storage capacity should be assigned to the system storage pool.
4. Minimum size of a system volume is 2 GB.

Steps:

Perform the following steps to determine storage capacity:

1. Determine the capacity and compatibility of existing storage against the required storage.

   Note: Consider temporary excess capacity that might be required to execute the selected data migration strategy.
2. Determine how much additional storage is required.

Related topics:

- "Determining the SAN physical configuration"
- "Determining the data migration process" on page 21
- "Perform capacity planning"

Determining the SAN physical configuration

Steps: Perform the following steps to determine the SAN File System physical configuration:

1. Determine if you need more than one SAN File System file system instance.

   For example, use two file system instances when you need to keep data from two departments completely separated. Separate SAN File System server

storage pools and client access filesets might be required for each instance. Each instance of SAN File System requires a master and at least one subordinate Metadata server.

2. Specify the metadata and data-block storage capacity requirements for each file system instance.
3. Define the availability (connectivity) requirements for each file system.

Related topics:

- "Determining the cluster configuration" on page 14
- "Determine SAN File System client configuration" on page 15
- "Determine SAN attachments"
- "Determining SAN File System LUN configuration" on page 16
- "Determining policy-based management structure" on page 19
- "Determining a security model" on page 20
- "Determining the data migration process" on page 21
- "Creating an installation plan" on page 22
- "Creating a site backup strategy" on page 23

Determine SAN attachments

Steps:

Perform the following steps to determine the SAN attachments:

- Determine how the engines hosting the SAN File System server cluster will attach to the SAN
- Determine the SAN zoning strategy

Related topics:

- "Determine Metadata server cluster zoning and SAN configuration"

Determine Metadata server cluster zoning and SAN configuration

Steps: Perform the following steps to determine how the engines hosting the cluster will attach to the SAN:

1. Determine the number of clusters that are needed to manage the file systems and provide the required metadata performance bandwidth.

   The number of engines required for the installation is directly related to the volume of transaction requests from the servers. To determine the number of engines required for the installation, do the following:

   a. Isolate the name space that is associated with the set of applications that will be using the SAN File System as a file system.
   b. Partition the name space based on the expected workloads. Partition the name space for disjoint workload sets that are independent of each other using a logical directory hierarchy.
   c. Specify which partition will be assigned to which SAN File System fileset.
   d. Determine the number of client machines (SAN File System clients). These machines are also the application server machines
   e. Specify which partitions from (step 1b) will use which machines. Consider the locations, typical application loads, and so forth. Consider the worst-case workload rather than the average workload.

f. Assume that each SAN File System client sends only 30% of its file system requests to the metadata server, and that the remaining 70% of file system operations are serviced from the client's local cache.

g. Determine the expected Metadata server transaction rates that will begin from each client per partition.

h. Determine the number of engines that will be needed by dividing the expected transaction rates from step 1g by the typical Metadata server transaction rates.

Note: Dynamic name space could be handled separately and spread over all servers or within the same partitioning scheme.

2. Switch zone (hard zoning preferred) each of the Metadata server host bus adapters (HBAs) separately such that each HBA can detect all storage, both Metadata storage and storage pool storage, but no HBA can see any other HBA. This means each HBA will reside in a separate zone.
3. LUN masking must be used where supported by the Metadata storage subsystem to LUN mask the Metadata Storage LUNs for exclusive use by the Metadata servers.
4. Specify that the Metadata server storage or LUNs are to be configured to the Linux mode (if the Metadata Storage Subsystem has OS-specific operating modes).
5. Install a SAN manager (platform and software), if needed. Note that Tivoli SAN Manager is available on the master console.
6. Develop the needed SAN topology roadmap to meet the customer requirements. Due to the limited number of LUNs addressable by the Metadata servers and the need for pathing redundancy, some combination of zoning or physical switch attachment should be used to ensure that each Metadata server has only two paths to all storage. One way to accomplish strict Metadata server dual-pathing is to use two separate fabrics, each with only a single path to storage. This provides fabric redundancy and prevents any switch from being a single point of failure. Another way to achieve strict Metadata server dual pathing is through the use of zoning. Zoning can be used to restrict the number of paths from each Metadata server to one path from each HBA port to the storage, even if the fabric connectivity itself allows many more paths to be found.
7. Determine what SAN HBA, switches, cabling modifications, upgrades, replacements, or additions are required.
8. Determine the SAN software modifications required. IBM Subsystem Device Driver (SDD) v1.4 is required for SANs using IBM Enterprise Storage Server or SAN Volume Controller.

Related topics:

- "Determine client and overall zoning and SAN configuration"
- "Determine SAN attachments" on page 12
- "Determine client and overall zoning and SAN configuration"

Determine client and overall zoning and SAN configuration

Context:

When planning SAN File System zoning strategy, issues such as the following must be considered: keeping metadata storage away from access by other systems on the SAN, the number of addressable LUNs, SAN security, and SAN component interoperability.

The SAN should also be appropriately zoned to prevent unauthorized access.

Steps:

Perform the following steps to configure SAN zones:

Note: To comply with the current restriction on the number of LUNs that the Metadata servers can access, you can attach the SAN by constructing two separate fabrics to control the number of paths to the storage. Each of these fabrics should be constructed of one or more switches from the same vendor. Constructing a single-path attachment from each HBA port of each Metadata server through redundant fabrics (achievable by zoning) can limit the number of paths found and maximize the number of actual physical LUNs which may be used with SAN File System, while still providing the necessary redundancy in case of path failure.

1. Create a zone for each Metadata server HBA that encompasses the server and its single HBA and all storage – both metadata storage and user pool storage. Keep in mind that there is no level of zoning you can do on a SAN that will protect SAN File System systems from SAN events caused by other non-SAN File System systems connected to the same fabric. You should not create fabrics that include traffic and administrative contact from non-SAN File System systems.
2. Create a zone for each HBA resident in each SAN File System client, together with the user data volumes to which you want that client to have access.

Related topics:

- Chapter 2, "Preparing the SAN", on page 29
- "Determine Metadata server cluster zoning and SAN configuration" on page 12

Determining the cluster configuration

Context: The SAN File System Metadata server is built to run in a clustered environment. It is hosted on a cluster system of from two to eight engines. The engines in the cluster communicate with each other over an IP local area network (LAN). An administrative server is installed on each engine. One administrative server in the cluster is appointed as the *master*. After a cluster is constructed, master status can be moved to another server if desired, or the situation requires it. The master is responsible for load balancing and physical space allocation. The unit of load balancing is the fileset, with the master engine assigning work to others based on the static load map. This static fileset:server binding is defined when the fileset is created. All engines in the server cluster are required to be homogeneous.

The SAN File System architecture requires SAN File System clients to be able to share access to LUNs with the cluster. Depending on the configuration (that is, the use of switch zoning), multiple SAN File System clients need to share access to LUNs. The SAN File System clients may be a mix of all types of client OSs supported in this release of SAN File System (that is, heterogeneous clients).

Steps:

Perform the following steps to determine the cluster configuration:

1. Determine the hardware configuration of each server cluster, which includes determining the number of server engines that are used in each cluster and their level of SAN/LAN connectivity.
2. Determine the location of the SAN File System console.

Related topics:

- "Determine SAN File System client configuration"

Determine SAN File System client configuration

Steps:

Perform the following steps to plan for individual SAN File System client considerations:

1. Determine if there is SAN File System client support for the client system's operating system type and version. This includes determining if an operating system upgrade to a SAN File System client is needed.
2. Determine if the SAN File System clients are compatible with SAN HBA drivers and volume management software.

   Note: You must isolate volume management software from SAN File System-managed storage devices. Volume management software may only be used on LUNs that reside outside the SAN File System storage pool.
3. Determine if there is adequate and compatible hardware and software to support the SAN File System file system availability and performance requirements.

   Determine if you need to add or replace SAN host bus adapters or add network interface cards (NICs) to meet the connectivity requirements.

   Determine if the firmware installed on these host bus adapters needs to be upgraded to meet SAN File System requirements.

   Determine if software changes such as installation and configuration of SAN multipath software, such as IBM Subsystem Device Driver (SDD) v1.4 are required.
4. Determine which data or storage devices will be redeployed within the SAN File System server storage.
5. Determine which applications (including backup applications) will be affected by SAN File System data migration or storage redeployment. You must decide if you need to reconfigure, correct, or replace the application
6. Determine if new client applications are compatible with SAN File System storage.

   Note: Any applications for which data has been relocated may require reconfiguration to target its new data location. Due to dependencies upon a particular file system, some applications may be incompatible with the SAN File System file system type. The incompatible applications must be repaired, replaced, or withdrawn. For example, some Backup products have file system type dependencies and, therefore, are unlikely to be compatible with SAN File System. Other applications may be affected by the introduction of sharing data with other clients' systems.

Determining SAN File System LUN configuration

You must determine how the newly-configured network will be structured and organized. Use the following picture for reference.

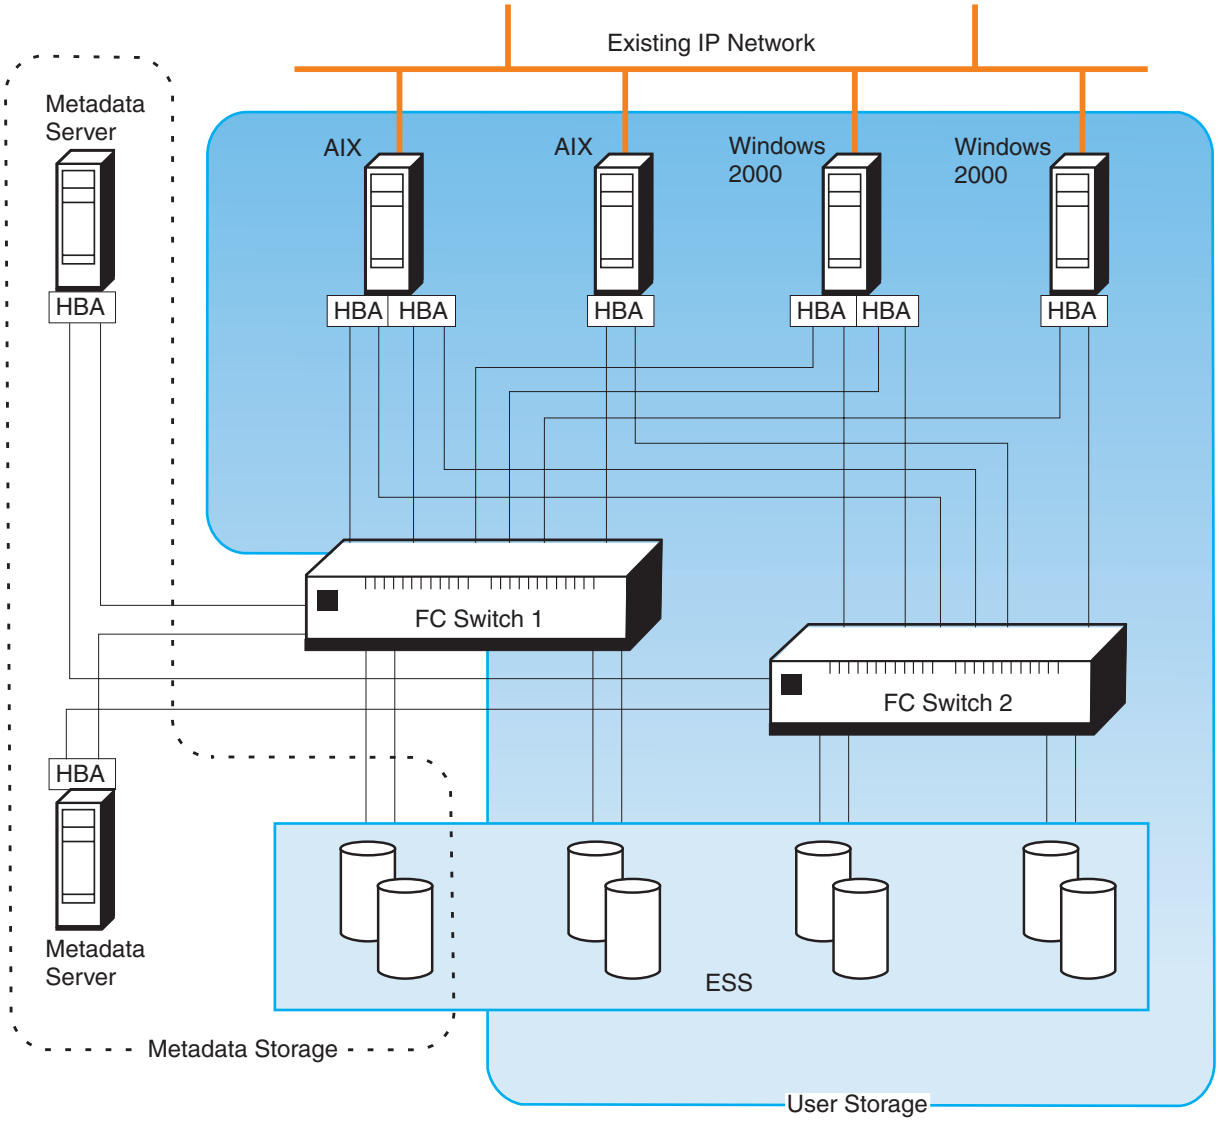

Figure 1. Example of a basic SAN File System

Related topics:

- "Determining which LUNs can be used as SAN File System volumes" on page 17
- "Determining which volumes will be used as metadata storage" on page 18
- "Determining the user storage pools needed" on page 18
- "Determining the filesets that will be needed" on page 19

Configuring LUNs

LUN is the *logical unit number* of a device.

Prerequisites:

As part of SAN File System configuration, logical unit numbers (LUNs) must be generated on the storage subsystem and mapped to the clients and Metadata servers. Each storage or disk partition in the subsystem has its own LUN. If the Metadata server wants to access the storage, it must request access to the LUN in the storage device.

The storage pool LUNs must be shared among all Metadata servers and all clients attached to SAN File System to ensure that the same LUNS are seen by all Metadata servers and clients. The Metadata server storage LUNs must be shared among all Metadata servers only. For many storage subsystems, LUN masking will be required to achieve metadata storage LUN isolation.

Steps:

Perform the following steps to configure LUNS for SAN File System:

• "Configuring storage pools" on page 57

Note: Certain events in the operating environment, such as operating system reboots and cable disconnects can cause SAN File System to lose connectivity to the LUNs. If there are any I/O failure errors logged on the client or server, you should ensure the following:

- Configured LUNs are visible from the SAN File System servers and the SAN File System clients. From the server, you should see both the user space LUNs and system LUNs. From clients, you should see all the user LUNs and no system LUNs.
- The fabric switch has not lost the zoning configuration. When the operating system is rebooted, it is possible for the fabric to lose the zoning configuration, thus preventing the nodes from reaching targets.

If connectivity has been lost, it might be necessary to force the SAN File System server and clients to remap the LUNs. Contact your SAN administrator for assistance with LUN rediscovery options specific to your operating environment.

Related topics:

• "Configuring storage pools" on page 57

Determining which LUNs can be used as SAN File System volumes

Steps:

Perform the following steps to determine which LUNs to use:

1. Define the LUNs to be used as SAN File System volumes taking into account the Linux LUN limitation.

   Keep in mind when mapping LUNs in Linux that there will be a system hard drive present, which could be the boot drive.

   The minimum size for system volumes is 2 GB.

   Notes:

   a. Ensure that you know what disk you are accessing before mapping LUNS; you could delete the boot drive.

   b. There is a strict limit of major/minor numbers available for SAN File System LUNs on the Metadata servers. See the Release Note for details. Multipathing decreases the number of available LUNs. Therefore it is

recommended that the SAN be configured to present only two paths from the storage to the Metadata server (one path per MDS HBA port). This provides addressability to the largest number of redundantly pathed LUNs (dual-pathed) possible. You can do this either by the use of physical attachment or through the use of hard zoning.

2. Specify that each LUN must be configured to be in heterogeneous access mode (that is, not set to a specific operating system type). For IBM TotalStorage Enterprise Storage Server, use ESS Specialist to set the modes.

   Note: The only exception is for Metadata server storage. The Metadata servers' LUNs should be configured for access only by the Linux operating system.

3. Ensure that all SAN File System storage uses LUN masking where possible to prevent disk access by non-SAN File System hosts or clients. For ESS, LUN masking is essential to prevent access to metadata storage from client systems.

Related topics:

- "Determining which volumes will be used as metadata storage"
- "Determining the user storage pools needed"
- "Determining the filesets that will be needed" on page 19
- "Determining SAN File System LUN configuration" on page 16

Determining which volumes will be used as metadata storage

Context: This information is required when you run setupTank to start the Metadata servers. You will need a list of IP addresses of subordinate nodes.

The SAN File System metadata must be configured to be in a separate storage pool and, therefore, on separate volumes than the user data. The metadata volumes can be on the same storage subsystem as the data volumes, or they can be on a different storage subsystem. Putting the metadata and data on separate storage subsystems has an impact on LUN-based backup and restore operations, but is a valid SAN configuration. The metadata volumes must be zoned away from the SAN File System client systems to avoid any possibility of metadata corruption by client systems or HBA interaction between client systems. This means that the switch zones must be set up such that the Metadata servers can detect all storage, but the client systems are zoned such that they can only detect the user storage pool. However, for some storage subsystems which allow access to all internal LUNs through any port (such as ESS), zoning is insufficient to provide this metadata storage isolation. For these subsystems, LUN masking is mandatory and must be used along with zoning to ensure metadata integrity. You must use LUN masking to reinforce LUN access integrity where supported by the storage subsystem being used.

Volumes have names that must be unique within a cluster, and can optionally have administrative descriptions.

Determining the user storage pools needed

Context: Storage pools can be organized based on any criterion that an organization chooses. One criterion could be device capabilities. Therefore, the degree of availability or the level of performance might be the primary organizational dimension. Other examples include business divisions, applications, business owners, localities, and customers.

Because the metadata grows as the file system grows, the storage administrator must be sure to assign enough volumes to the system pool. The system pool requires approximately 10% of the total storage capacity that SAN File System manages.

The files in a fileset can belong to different storage pools. Multiple filesets can own storage within a single storage pool.

As part of the planning process, you should determine which types of storage pools are needed.

Keep in mind that object names for user pools must be in ASCII only.

Determining the filesets that will be needed

Context:

The goal in SAN File System is to make all engines busy in a balanced manner. This is facilitated through the use of filesets.

You must have at least *N* filesets in an *N* engine cluster or engines will unintentionally be in standby mode. This is because a fileset can be managed by one engine only. You should carve out the workload into a multiple of *N* filesets, all expected to be similar in terms of server workload.

Note: Because engine workload is a function of metadata traffic and not data traffic, the notion of workload can be deceiving. For example, a database server with several terabytes of data in a warehouse doing few writes of new files or added blocks might be less busy in terms of server traffic than a mail server for a few users that intermittently creates hundreds of pieces of spam per user per day.

Determining policy-based management structure

Context:

SAN File System includes powerful mechanisms for controlling how administrators manage the files in the global file system. SAN File System policies primarily involve the placement of files in storage pools using rules based on file attributes such as file name, owner, group ID of owner, or the system creating the file.

Some sample policy sets are provided for you in Appendix G, “Sample policy sets” , on page 117.

Steps:

Perform the following steps to determine a policy-based management structure:

1. Determine the placement rules for the files.
2. Determine the policies.

   Notes:

   a. SAN File System can have multiple policies defined within its configuration, but only one of them at any time is actively being used.
   b. A default policy is provided with SAN File System, called DEFAULT_POLICY.

c. Placement based on creation time, user ID, or group is not reccomended because none of these attributes work correctly when files are being restored from backup, or being migrated. In such cases, the creation time is always the time of the restore or migration, and the user and group are always those of the restore or migration application.

Related topics:

- Appendix G, "Sample policy sets", on page 117
- "Creating a policy" on page 65

Determining a security model

Context: When determining the security model, keep in mind SAN File System provides heterogeneous access control for the files in the file system, which means:

- When files are created and accessed from Windows clients, all the security features of the Windows platform (this includes DAC, implemented using Windows ACL) are available and enforced.
- When files are created and accessed from UNIX clients, all the security features of UNIX platform (read/write/execute permissions for owner, group and others) are available and enforced.
- When files created in a UNIX client are accessed in a Windows client, the access is controlled only using the semantics (and the permissions) of *other* in UNIX . Similarly, when files created in a Windows client are accessed in a UNIX client, the access is controlled only using the semantics (and the permissions) of *Everyone* in Windows. Individual user and group access control is not possible for files when files are accessed across domains (UNIX and Windows clients).

Steps:

Perform the following steps to define a security model:

1. Set up the IP network so that all systems hosting Web browsers to be used for SAN File System administration have IP access to the SAN File System systems hosting the Administrative servers and Metadata servers (the engines).
2. Configure the environment to define the users who can administer SAN File System and the administrative roles that are assigned to those users.
3. Ensure that all clients for AIX share a common user name space, either through NIS or compatible /etc/password file. Ensure that all clients for Windows share the name space through a Windows directory service.

Security model for defining systems for SAN File System administration

Context: When planning for a SAN File System configuration, you must determine which systems will be used for SAN File System administration and set up the customer's IP network so that all systems hosting Web browsers to be used for SAN File System administration have IP access to the SAN File System systems hosting the Administrative servers and Metadata servers.

Related topics:

- "Determining a security model"
- "Defining SAN File System users and roles" on page 21
- "Security" on page 33

Defining SAN File System users and roles

Context: During the planning phase, you must determine the users who can administer the SAN File System and the administrative roles assigned to those users. Users and their assigned roles are defined in the LDAP server used by the SAN File System cluster. After authentication, roles are based on the group assigned to the user during the login process.

SAN File System supports the following roles:

- Monitor
- Backup
- Operator
- Administrator

Determining the data migration process

Data migration is an optional part of installation.

Having a migration strategy is vital to move the data safely to SAN File System. Currently, the data migration process migrates data from an existing file system to a SAN File System file system at file level by using a user-level application that copies regular files, directories and symbolic links. Each file object is simply copied from the original file system to the target file system. This is performed using the SAN File System migration tool, which is a command line-based application.

There are several items that must be determined before attempting to migrate the data:

- The source file system cannot be in use by user applications during the migration. If needed, data can be migrated while user applications are updating the source file system. However, this makes verification difficult after the migration process is complete.
- The person running the migration tool must have enough privileges to copy the data and preserve the file system object attributes on the SAN File System file system.
- The amount of storage space required during the migration is twice that of the original system. The amount of storage space required in the destination SAN File System file system is always slightly more than the space currently consumed in the source file system. For example, if the data in the source file system occupies *n* blocks, during migration the SAN File System file system allocates at least *n* blocks. Therefore, the total storage space required is more than 2**n* blocks.
- The migration tool must be run by a privileged user so that the attributes of a file system object can be preserved after migration.

Migration process steps:

The migration tool is designed to handle the migration in three steps:

1. Gather environment data

   This step determines the resources available (RAM, disk space, and so on) for the migration, and estimates the time required to complete the job.
2. Create placement policies

   This step prepares the Metadata server cluster for the addition of new files.
3. Migrate the data

This step consists of the actual data migration. In this step the files are copied from the source to the destination SAN File System file system.

4. Verify

   This phase verifies that the data was copied correctly. Data verification actually happens during the migration phase, but it can occur separately after the migration phase to ensure the consistency of the data.

Related topics:

- "Placement policies" on page 59
- "Migrating data" on page 71
- "Data-migration prerequisites" on page 6

Creating an installation plan

Steps:

Perform the following steps to create an installation plan:

Note: User applications should be installed on local drives. For more information, see Appendix H, "Managing local drives", on page 119.

1. Determine the pre-SAN File System installation, which includes:
   - Hardware modifications
   - Software modifications
   - Additional storage
   - Data relocation

   Note: Fibre channel and Ethernet cables do not ship the SAN File System hardware. They must be provided by the customer.
2. Specify the hardware installation of:
   - SAN File System server engines
   - Local Area Network (LAN) cabling and other cabling to clients
   - SAN switch insertion
   - Cable modifications
   - Zone reconfiguration to server-only storage
   - Zone reconfiguration of data-block storage shared with clients
3. Determine the buildup schedule of the target configuration for SAN File System servers and clients. Schedule the:
   - Client software installation and reconfiguration
   - Data migration
   - Addition to or redeployment of existing storage into SAN File System

     Note: Space can be deployed as server metadata storage or as aggregated file system data storage. Reconfiguration of servers, clients, SAN topology, and storage devices might be necessary. You must plan which drives will be metadata disks. Also, you will need a list of the IP addresses of subordinate nodes.
   - Migration or reconfiguration of affected applications
4. Determine the personnel resources required for installation.
5. Validate hardware and software delivery dates.

6. Schedule the installation process.

Related topics:

- "Perform site audit" on page 6
- "Perform capacity planning" on page 11
- "Determine SAN attachments" on page 12
- "Determining SAN File System LUN configuration" on page 16
- "Determining a security model" on page 20
- "Determining the data migration process" on page 21

Creating a site backup strategy

Steps:

Perform the following steps to create a backup strategy:

1. Modify the backup strategy for the data that is relocated within clients but is not part of SAN File System.
2. Determine the strategy for client backup and restoration of visible filesets.
   a. Determine which client is responsible for backing up a shared file system. Files and directories created on a Windows operating system should be backed up and restored through Windows, and UNIX® files and directories should be backed up and restored through the UNIX operating system.
   b. Determine which backup applications are compatible with the SAN File System file system.
3. Determine the backup and recovery strategy for SAN File System servers and which application to use.
4. Determine if new, additional, or upgraded backup devices are required because of performance, SAN File System compatibility, or modified application compatibility considerations.
5. Determine the SAN hardware modifications in addition to zone and operating system configuration procedures for sharing backup devices among systems, if required.

Related topics:

- "Backup and restore"

Backup and restore

Backup is the process of saving copies of your files, and *recovery* is the process of restoring those copies if your original files are damaged or lost. For SAN File System, an administrator must also back up the system metadata, which includes information about fileset attachment points, storage pools, volumes, and file placement policies. This backup data is used to re-create cluster configuratoin if necessary.

Although SAN File System does not provide specific backup-and-restore operations for files, it 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), SAN File System clients can use the functions and features of that product to back up and restore files that reside in the SAN File System global namespace.

To perform LUN-based backups, an administrator can use the copy services features that exist in the storage subsystems that SAN File System supports.

To create a backup copy of the system metadata configuration, an administrator can use the Disaster Recovery task on the SAN File System console or an administrative command.

Using a file-based approach:

To back up and restore files in the global namespace using a file-based approach, users can run standard tools and utilities on SAN File System clients.

The first line of defense for scenarios where files have been lost but the overall system remains healthy is the use of the FlashCopy image function in SAN File System. To assist with the backup process, an administrator can choose to create FlashCopy images of filesets that can be backed up at a later time. A FlashCopy image contains read-only copies of the files in a fileset as they exist at a specific point in time.

The FlashCopy image is stored in a special subdirectory named .flashcopy under the fileset's root attachment point. After an administrator creates a FlashCopy image of a fileset, a user can use standard backup tools to back up the files from a SAN File System client by specifying the path to the FlashCopy image instead of the path to the actual files.

Note: Individual files could be copied out of the .flashcopy directory if less than the entire fileset needs restoration.

Users and applications can continue working with the actual files while the backup occurs.

Understanding restrictions for a file-based approach:

Both backup administrators and users on client machines must be aware of restrictions that apply if files are backed up for use on both AIX® and Windows® clients.

First, an AIX user who requests the backups must have read permissions on all files and search permissions on all directories. This is typically a root user. A Windows user must have read permissions on all files and "list folder contents" permissions on all folders.

Second, to avoid losing security attributes, users must organize the file system in a special way. There are two options:

- The first is to ensure that the root of each fileset contains only directories and folders, and use a naming convention for each that makes it clear to users whether a directory contains files created by an AIX client or by a Windows client. The top-level directories can contain any combination of files, links, and directories, just as they can within NFS and CIFS.
- The second is to choose either AIX or Windows as the global namespace default, and to require (by convention) that files, directories, and links created by non-default clients be created in specially named directories. Users can place these specially named directories at any level in the file system; however, any file, link, or directory beneath them must also be created by the same type of client.

In either case, the special naming conventions are for user and administrator benefit only. SAN File System does not interpret them in any way. In addition, if a directory or folder grants search or "list folder contents" permissions to "Other" or "Everyone," SAN File System does not prevent an administrator or any client user from creating a file that violates the convention.

Therefore, administrators and users must be aware of the naming conventions when performing backups, and ensure that backups of files in specially named directories are performed only from the same type of client that created the files. This means that, in an environment that has both AIX and Windows clients, the backup process must be divided into multiple parts to prevent the loss of security attributes for files. The number of parts could be as few as two, or as many as the number of directories in the global namespace, depending on the capabilities of the native backup utility a client uses.

Sometimes disaster strikes and the system — as well as the FlashCopy images— are unusable. For these disaster recovery scenarios, there are two basic approaches for backup and restoration:

1. The LUN method
2. The API method

While these methods are not mutually exclusive, your disaster recovery plans will primarily involve one or the other.

The LUN method saves and restores data at the device level (that is, a JABOB – Just a Bunch of Bytes – approach); the API method saves and restores data at the file level. For a variety of reasons, the LUN method is simpler to manage, and provides more coverage than the API method. To adopt the LUN method, however, the actual copying and restoring of data must be provided as a service by the underlying storage subsystem. If your storage subsystem meets this requirement, the LUN method is the recommended backup and restore approach. If your storage subsystem provides no such service, the API method is your only available option.

Note: This service does not have to be centralized and homogenous, covering the entire SAN, although such a service simplifies the procedure. You may choose to pursue the LUN method even for a fragmented SAN that requires a piecemeal LUN copy across two or more storage subsystems. In such a scenario, you would be responsible for manually managing those multiple backup sets as though they were a single backup set.

Using the LUN-based approach:

To back up the entire SAN File System global namespace in a single operation and to restore the global namespace as a complete namespace, an administrator can use a LUN-based approach. The administrator can use the copy services features that exist in the storage subsystems that SAN File System supports, such as the FlashCopy feature of the IBM TotalStorage Enterprise Storage Server.

When performing a LUN-based backup, an administrator must be sure to back up both the LUNs used as volumes in user storage pools and the LUNs used as volumes in the system storage pool (which is used for metadata) at the same time.

Before performing a LUN-based backup, an administrator must quiesce the Metadata server cluster. This is required for a consistent backup. An administrator

can also choose to stop the cluster (from the SAN File System console or by using the **stopcluster** command) before performing a LUN-based backup.

Using the API-based approach:

Considerations for filesets within directories for the API method.:

There are special API method considerations when filesets reside within directories. From the SAN File System client perspective (and therefore from the backup application perspective), a fileset within a directory looks exactly like a regular subdirectory. From the Metadata server and Administrative server perspective, however, this is a fileset that was attached to an arbitrary subdirectory in the global namespace.

The problem is that backup applications generally create subdirectories on-the-fly when attempting to restore files to a directory that does not exist. In the case of SAN File System, the subdirectory being created might have originally been part of another directory. This problem is avoided by recreating all subdirectory connections prior to restoring all the regular subdirectories and files.

The **mkdrfile** command saves state information that can be used to recreate this portion of the global namespace following some disastrous loss of data. Unfortunately, regular subdirectories can only be created from a client, therefore, the Metadata server cannot recreate directory trees which contain a mix of filesets within subdirectories, and regular subdirectories. To avoid this problem during disaster recovery scenarios, it is strongly recommended that you only attach filesets to the cluster root directory (usually seen as /mnt/sanfs from the client), and to each other. In other words, as a best practice, do not attach filesets to regular subdirectories. The **mkdrfile** output can then be used to completely restore the top of the global namespace tree before using the client-based backup application to restore the rest of the global namespace.

FlashCopy image considerations for the API method:

The .flashcopy subdirectories created when FlashCopy images are made, are read-only. Unless told otherwise, the client backup application would typically backup the .flashcopy subdirectories along with everything else. At restore time, however, the same backup application would attempt, but not be able, to redeposit the original files into these subdirectories. (You would not want them there anyhow since they would appear to be valid FlashCopy images from the client perspective, when in reality, the metadata needed for the original FlashCopy images would be missing).

One obvious consequence of this behavior is that there is no way, using the API method, to restore your original FlashCopy images if you have lost your metadata in a disaster scenario. You only get the original files restored. However, if your backup application supports the ability to restore files to a directory other than their original location (that is, to the grandparent directory two levels above the ./.flashcopy/<flashcopyname> directory), then you have all the ingredients for a highly efficient API method backup, which leverages the FlashCopy image feature.

In any case, periodic FlashCopy images are still highly recommended. They are the most efficient method for quickly backing up and restoring files in scenarios where the metadata is still available.

Backing up and restoring system metadata:

SAN File System manages file data and system metadata separately. When a user backs up a file, only the file's data and attributes are backed up. For disaster recovery purposes, an administrator must back up system metadata (which includes information about fileset attachment points, storage pools, volumes, and policies), separately.

An administrator can create a file that contains a backup copy of system metadata from the SAN File System console or by using the **mkdrfile** command. The file, which is stored in the /usr/tank/server/DR directory on the master Metadata server's local disk, contains everything required to re-create the metadata. When needed, an administrator can use the contents of this file (along with normal restore processes for file data) to re-create the state of the cluster.

To restore system metadata, an administrator processes the information contained in the system metadata backup file using the **builddrscript** command. This command creates several scripts that the administrator must review to obtain a restore scenario, and then run to re-create the SAN File System configuration. Then, user data files can be restored from SAN File System clients.

An administrator should run the **mkdrfile** and **builddrscript** commands often enough to ensure that any configuration changes are reflected in the output of these commands. Note that an administrator should put copies of the output of the **mkdrfile** and **builddrscript** commands in an easily recoverable location on backup media where critical system and application files are kept for backup and restore purposes.

Note: To assist in protecting against the corruption of metadata and other metadata failures, an administrator can check metadata from the SAN File System console or by using the **startmetadatacheck** command. This command performs a consistency check on the system metadata, and optionally repairs any problems it finds. An administrator can check file metadata for one or more filesets, the system metadata, or both. There is also an option to check only the metadata structure, or to check the metadata structure and its contents.

There are three cases when an administrator might use the check or repair operation: 1) as part of a regularly scheduled cycle of preventive maintenance, 2) in response to an alert that recommends that this operation be performed (extra detail may be supplied that specifies the restore option that an administrator must use to salvage the metadata), and 3) if metadata corruption (or any other SAN File System corruption) is suspected. If the check metadata operation cannot resolve the problem, an administrator must perform a full restore of SAN File System, beginning with restoring the metadata. It is critical that **mkdrfile** is run in order to recover from such a situation.

Related topics:

- "Backup and restore planning"
- "Creating a site backup strategy" on page 23

Backup and restore planning

Steps:

While SAN File System does not provide a specialized backup and restore (BAR) solution, as with other file systems, traditional BAR methods and solutions still generally apply. However, there are some special requirements that should be considered when developing an off-the-shelf backup and restore solution for SAN File System.

Note: Some of these factors also affect SAN configuration choices, as some SAN configurations will be easier to back up than others.

The primary special requirement is the need to back up several disparate sets of data, then to collect and manage those sets of data as though they were a single, integrated backup set. The sets of SAN File System data that need to be backed up are:

1. The Metadata server disaster recovery file (generated by the **mkdrfile** command), which allows you to reconstruct filesets and their attach points. This file resides on the master Metadata server boot disk.
2. The Metadata server configuration files that define details such as cluster configuration, Administrative server configuration, and so forth. These files reside on the master Metadata server boot disk.

   Note: Some of the cluster configuration may alternately be recreated from the metadata disk known as the Master Volume.
3. The contents of the SAN File System metadata, which record where the client data is located on the SAN. This data resides in the server-only zone of the SAN.
4. The SAN File System client file data. This data resides in the open-to-all zone of the SAN.

Refer to Chapter 7, “Backing up”, on page 79 for recommended backup procedures for each of these SAN File System data components. The individual responsible for developing a SAN File System backup plan should be thoroughly familiar with these procedures prior to choosing and configuring an off-the-shelf BAR solution, and prior to finalizing decisions about how to configure the SAN for use by SAN File System. Proper consideration of these factors will lead to a more efficient and reliable SAN File System BAR procedure.

Related topics:

- “Backup and restore” on page 23
- Chapter 7, “Backing up”, on page 79
- “Creating a site backup strategy” on page 23
- Appendix F, “Disaster recovery”, on page 115

Chapter 2. Preparing the SAN

You must prepare your SAN for SAN File System installation.

Steps:

Perform the following steps to prepare the SAN for SAN File System installation:

1. Set up 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.
2. Create the SAN zones for SAN File System.
3. Resolve any hardware and software incompatibility issues.
4. If you have not yet established your placement policies, see "Placement policies" on page 59.

Related topics:

- "Configuring SAN zones" on page 30
- "Resolving incompatibility issues" on page 33

Setting up zones

When metadata storage and user storage reside on the same storage subsystem, special consideration must be made to 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. Metadata servers should be zoned or LUN Masked such that each can access all storage, both metadata and user storage. Client systems should be zoned or LUN Masked such each can see user storage only.

1. Switch zone (hard zoning preferred) each of the Metadata server host bus adapters (HBAs) separately such that each HBA can detect all storage, both Metadata storage and storage pool storage, but no HBA can see any other HBA. This means each HBA will reside in a separate zone.
2. LUN masking must be used where supported by the Metadata storage subsystem to LUN mask the Metadata Storage LUNs for exclusive use by the Metadata servers.
3. Specify that the Metadata server storage or LUNs are to be configured to the Linux mode (if the Metadata Storage Subsystem has OS-specific operating modes).

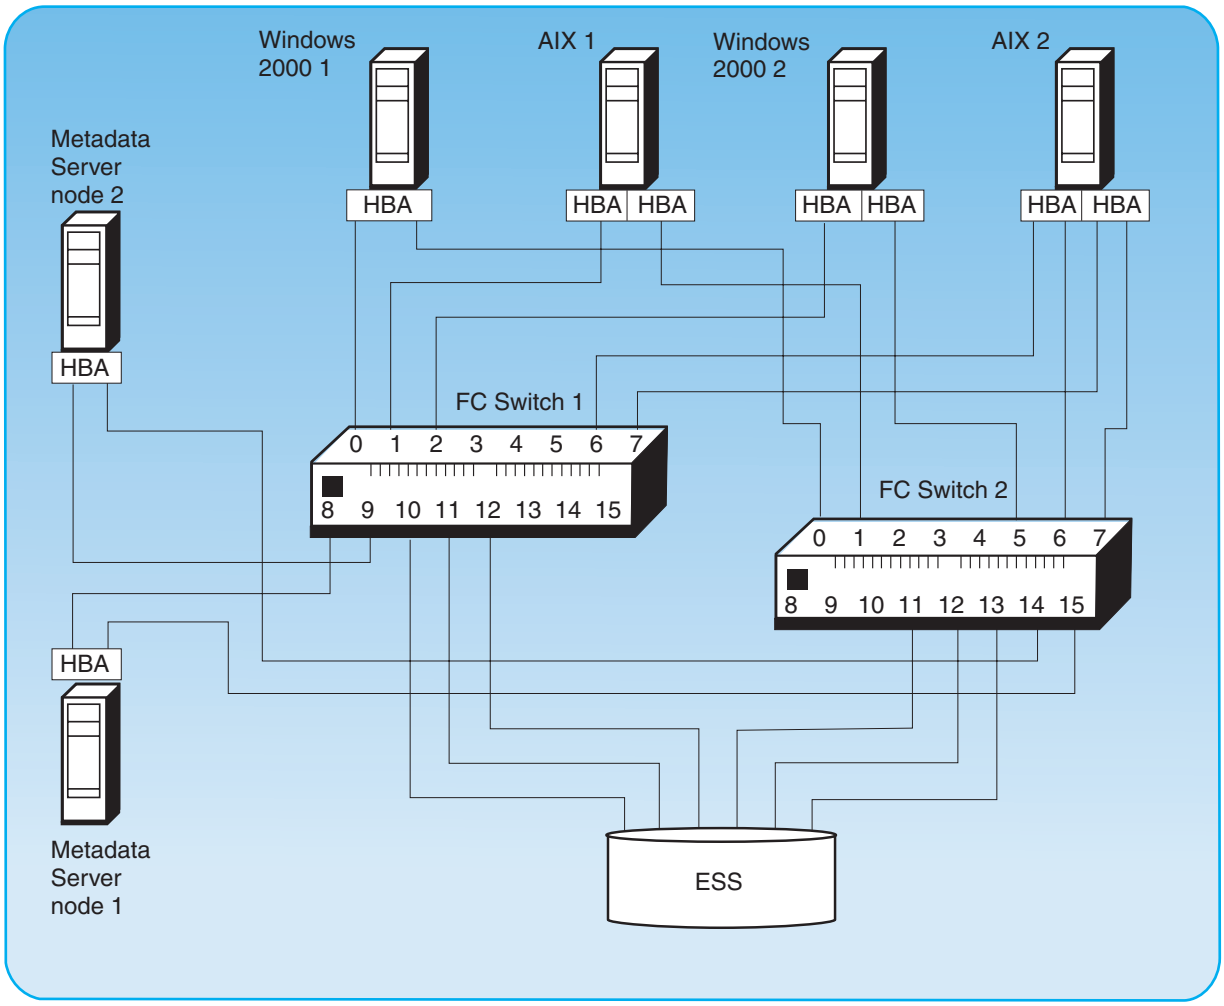

Figure 2. Example of a zoned network

Related topics:

- "Configuring SAN zones"
- "Determine client and overall zoning and SAN configuration" on page 13

Configuring SAN zones

Steps:

Perform the steps listed here to configure storage area network (SAN) zones:

Notes:

1. Due to the restriction on the number of LUNs the Metadata servers can currently access, special consideration must be made to limit of the number of paths created through the fabrics from each Metadata server to the storage to two (one per 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.
2. It is recommended the SAN File System user LUNs and SAN File System metadata LUNs do not share the same ESS 2105 Host Adapter ports.

3. Keep in mind that there is no level of zoning you can do on a SAN that will protect SAN File System systems from SAN events caused by other non-SAN File System systems connected to the same fabric. You should not create fabrics that include traffic and administrative contact from non-SAN File System systems.

When metadata storage and user storage reside on the same storage subsystem, special consideration must be made to 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. Metadata servers should be zoned or LUN Masked such that each can access all storage, both metadata and user storage. Client systems should be zoned or LUN Masked such each can see user storage only.

Example:

The following illustration shows an example of a network zoned for SAN File System.

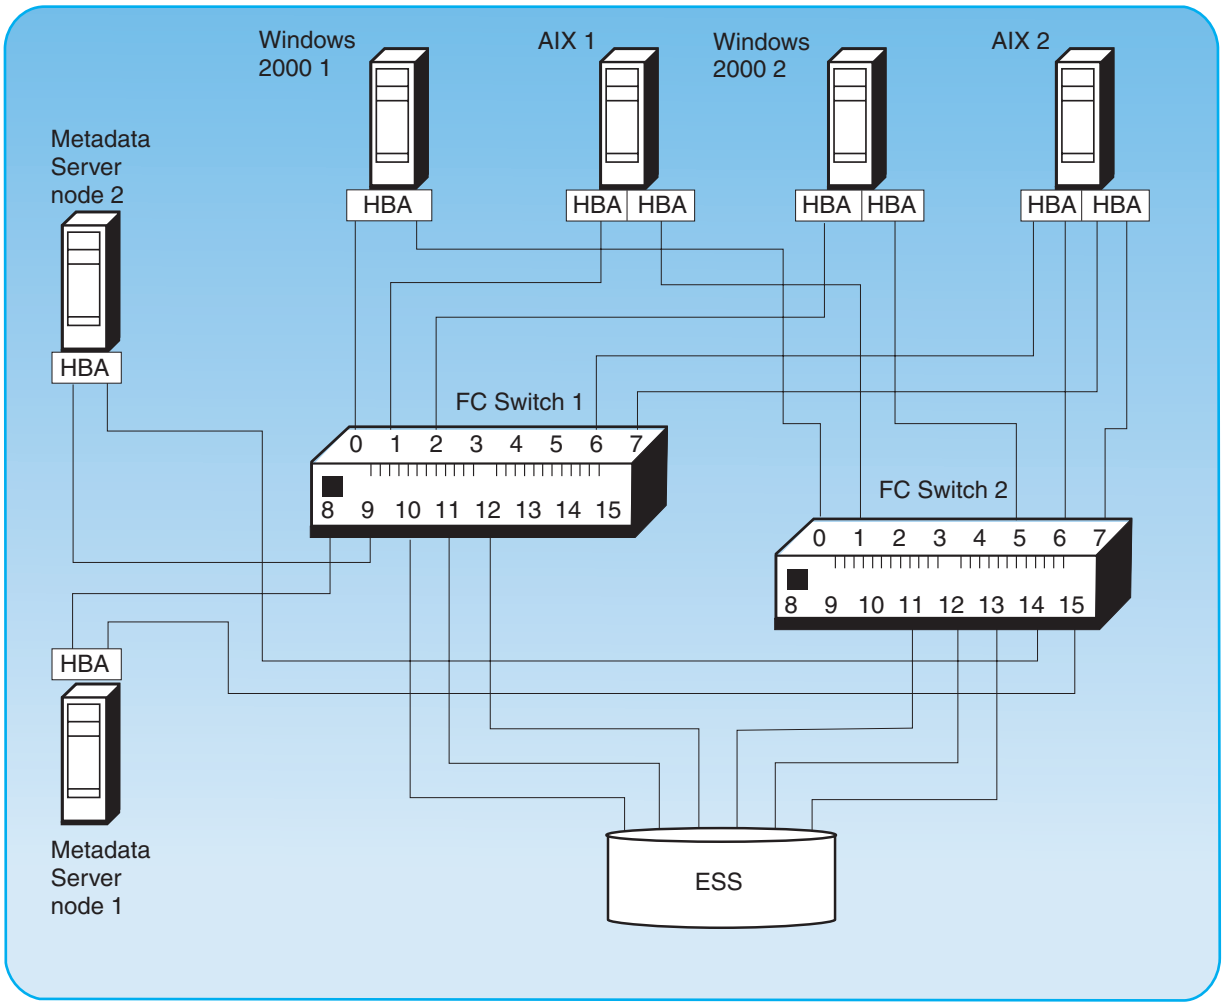

Figure 3. Example of a zoned network

Related topics:

- "Determine client and overall zoning and SAN configuration" on page 13
- "Setting up zones" on page 29

Reconfiguring Storage Devices

Prerequisites: The storage devices supported by SAN File System are:

- IBM TotalStorage Enterprise Storage Server, model 2105–F20
- IBM TotalStorage Enterprise Storage Server, model 2105–800
- IBM TotalStorage SAN Volume Controller, model 2145
- IBM TotalStorage SAN Integration Server, model 2146

Check the Release Notes for the supported code levels of these storage subsystems.

Context:

During installation and configuration of SAN File System, the ESS storage devices must be configured to work with the operating system used by the SAN File System servers. Use the documentation provided with your storage device to configure the device for generic use.

Related topics:

- "SAN File System planning worksheet" on page 102

Resolving incompatibility issues

Context:

Hardware and software that is incompatible with SAN File System must be isolated or resolved before installing SAN File System.

Steps:

To resolve incompatibility issues:

1. Isolate incompatible volume management software from SAN File System-managed LUNs. Volume management software can not be used to manage storage in the storage pool.
2. Isolate clustering packages including Windows 2000 MSCS and AIX HA-CMP from using storage that is in the SAN File System storage pool or in the SAN File System system disk pool.
3. Isolate all applications that use SCSI-2 Reserve or SCSI-3 Persistent Reserve protocols from accessing storage that resides within the storage pool.
4. Replace multipathing software that is supporting storage pool storage with the SAN File System-supported multipathing software.

Related topics:

- "Perform site audit" on page 6
- "Software prerequisites" on page 4

Security

SAN File System provides *administrative security* to protect against unauthorized access to SAN File System administrative operations, and *client security* to protect against unauthorized access to files in the SAN File System global namespace. For more information, see the related topics below.

Chapter 3. Physical Installation

Properties:

The following sections are included in physical installation. These are for the IBM service representative.

Note: The SAN File System hardware configuration should have been planned in advance by the customer. Determine what plans have been made before beginning the physical installation.

Unpacking the engine
This section describes how to remove the engine from the packing.

Installing the Model 1RX in a rack
This section directs you to the complete rack installation instructions that come with your Model 1RX engine.

System reliability considerations
This section provides information you should know to ensure reliable hardware operation.

Cabling
This section tells how to cable the SAN File System devices to each other and to the network.

Related topics:

- "Unpacking the engines" on page 36
- "Installing the Model 1RX in a rack" on page 36
- "System reliability considerations" on page 37
- "Cabling" on page 38

Installation and configuration checklist

Installation:

The Model 1RX engines come complete in the box. No devices must be installed in them.

1. Unpack the engines. See "Unpacking the engines" on page 36.
2. Install engines in the rack. See "Installing the Model 1RX in a rack" on page 36.
3. Cable the devices together. See "Cabling" on page 38.

Configuration:

The SAN File System software comes installed in each engine.

1. Install the master console. See "Installing the master console in a rack" on page 37.
2. Set administrative privileges. See "Setting administrative privileges and fileset permissions" on page 69.
3. Reconfigure the SAN storage devices. See "Reconfiguring Storage Devices" on page 32.

4. Perform LUN masking. See Appendix H, “Managing local drives”, on page 119.
5. Set up zones. See “Setting up zones” on page 29.
6. Reset IP addresses in the engines. See “Assigning IP addresses for Metadata servers” on page 52.
7. Reset the IP addresses in the RSA adapters. See “Setting RSA adapter IP addresses using the master console” on page 53.
8. Assign the master engine and create the cluster. See “Starting the Metadata servers with the setupTank utility” on page 55.
9. Install SDD on the clients. See “Installing Subsystem Device Driver v1.4 on clients” on page 66.
10. Install the clients. See “Installing client software” on page 66.
11. Additional configuration:
    a. Configure filesets. See “Configuring filesets” on page 58.
    b. Set up one or more policies. See “Placement policies” on page 59.
12. Backup and restore. See Chapter 7, “Backing up”, on page 79.

Unpacking the engines

This section is for the IBM service representative.

The following items are shipped with the SAN File System. Verify that you have all the items listed.

- From two to eight Model 1RX engines
- 2 Power cords per engine
- Documentation set, including:
  - *Safety Information—Read This First*
  - This document.
  - SAN File System Publications CD
  - Several additional CDs
- Rack-mounting kits, including:
  - Cable-management assembly
  - Cable-management-arm bracket
  - 2 Hinge pins
  - 5 Cable straps
  - 2 Cable-restraint brackets
  - 2 Slide rails
  - 5 Cable clamps
  - 5 Cable ties
  - 5 M6 screws

Related topics:

- “Safety information” on page 125

Installing the Model 1RX in a rack

This section is for the IBM service representative.

Complete instructions for installing the Model 1RX in a rack are included in the box with each. The rack installation document is part number 88p9186.

Notes:

1. Plan to allow space for the master console and Ethernet switches and hubs if they will be in the same rack.
2. Be sure to connect the cable management arm to the rail guide before installing the rail guide in the rack. After the rail guide has been installed, you may not be able to insert the pin that holds the cable management arm to the rail guide.

Related topics:

- "Safety information" on page 125

System reliability considerations

This section is for the IBM service representative.

To help ensure proper cooling and system reliability, make sure that:

- Each of the drive bays has either a drive or a filler panel installed.
- Each of the power-supply bays has a power supply installed.
- For rack configurations, make sure that space is available around the engine to enable the engine cooling system to work properly. See the documentation that comes with the rack for additional information.
- The engine cover is in place during normal operation.
- The air-baffle cover over the microprocessors remains closed during normal operation.
- The air baffle is installed between the fans and the power supply.
- A removed hot-swap drive is replaced within 10 minutes of removal.
- Cables for optional adapters are routed according to the instructions provided with the adapters.
- A failed fan is replaced within 48 hours.
- The engine is powered off and the power cords are disconnected before you remove an air baffle.
- The air baffle is always installed in the engine except when you are installing or removing the components that are located under the air baffle.

Installing the master console in a rack

Context:

Note: If there is a SAN Volume Controller or another SAN File System already installed on site, its master console may be used for the SAN File System as well. In that case, skip this section.

The master console comes fully loaded on an IBM @server xSeries 305. Follow the instructions shipped in the box to unpack the hardware and install it in a rack. Cable connections to the SAN File System and for the keyboard and display are described in "Cabling" on page 38.

Steps:

The following steps must be performed to complete the master console installation. After the hardware is installed and the cabling completed:

- Download and install the VNC server software on the master console.

- For remote access, establish a VPN tunnel using the Connection Manager utility.
- Configure the master console for Call-Home. See "Configuring Service Alert on the master console" on page 48.

Related topics:
- "Cabling"

Cabling

Steps:

For each engine perform the following steps to cable the hardware:

1. 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.
2. 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 6 on page 40 for a cabling example.
3. For redundancy, connect another Fibre Channel cable to the other HBA port in expansion slot 2, and to the other switch (or zone). This is optional, but recommended.

   Note: It is recommended the SAN File System user LUNs and SAN File System metadata LUNs do not share the same ESS 2105 Host Adapter ports.
4. 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 6 on page 40, shows a hub.
5. The RSA adapter comes preinstalled on each engine. An Advanced System Management (ASM) connector and USB cable are provided.
   a. Connect the USB cable to a USB port on the engine and the other end to the RSA card.
   b. 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 6 on page 40.
   c. Connect a ASM breakout cable (dongle) to the ASM connector on the RSA II card in each of the engines present in the SAN FS cluster. Connect the ASM breakout cable to the previous and next ASM breakout cables with RJ-45 cables. The first and last RJ-45 sockets in the chain must be terminated with the terminators provided. See Figure 4 on page 39. and Figure 5 on page 39.
   d. The 9-pin D-shell serial connector on the ASM conector is not used.

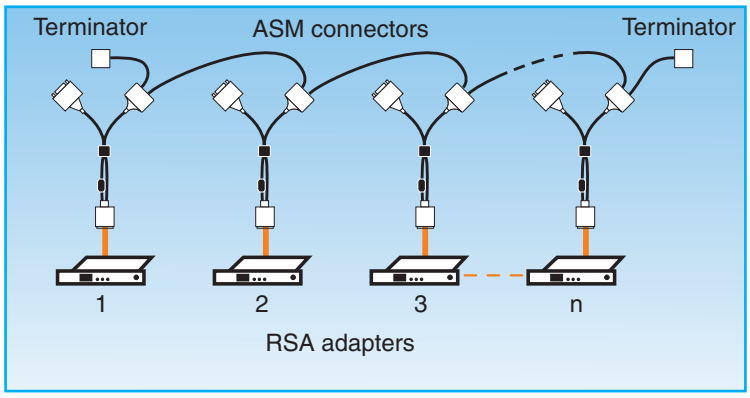

Figure 4. Connecting the RSA adapters together

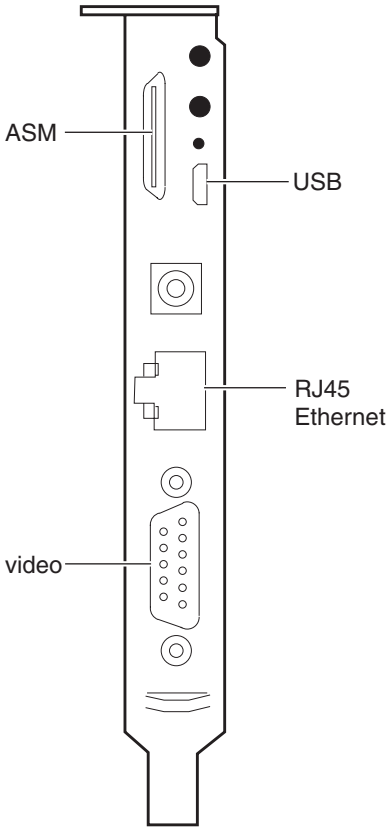

Figure 5. RSA II adapter connectors

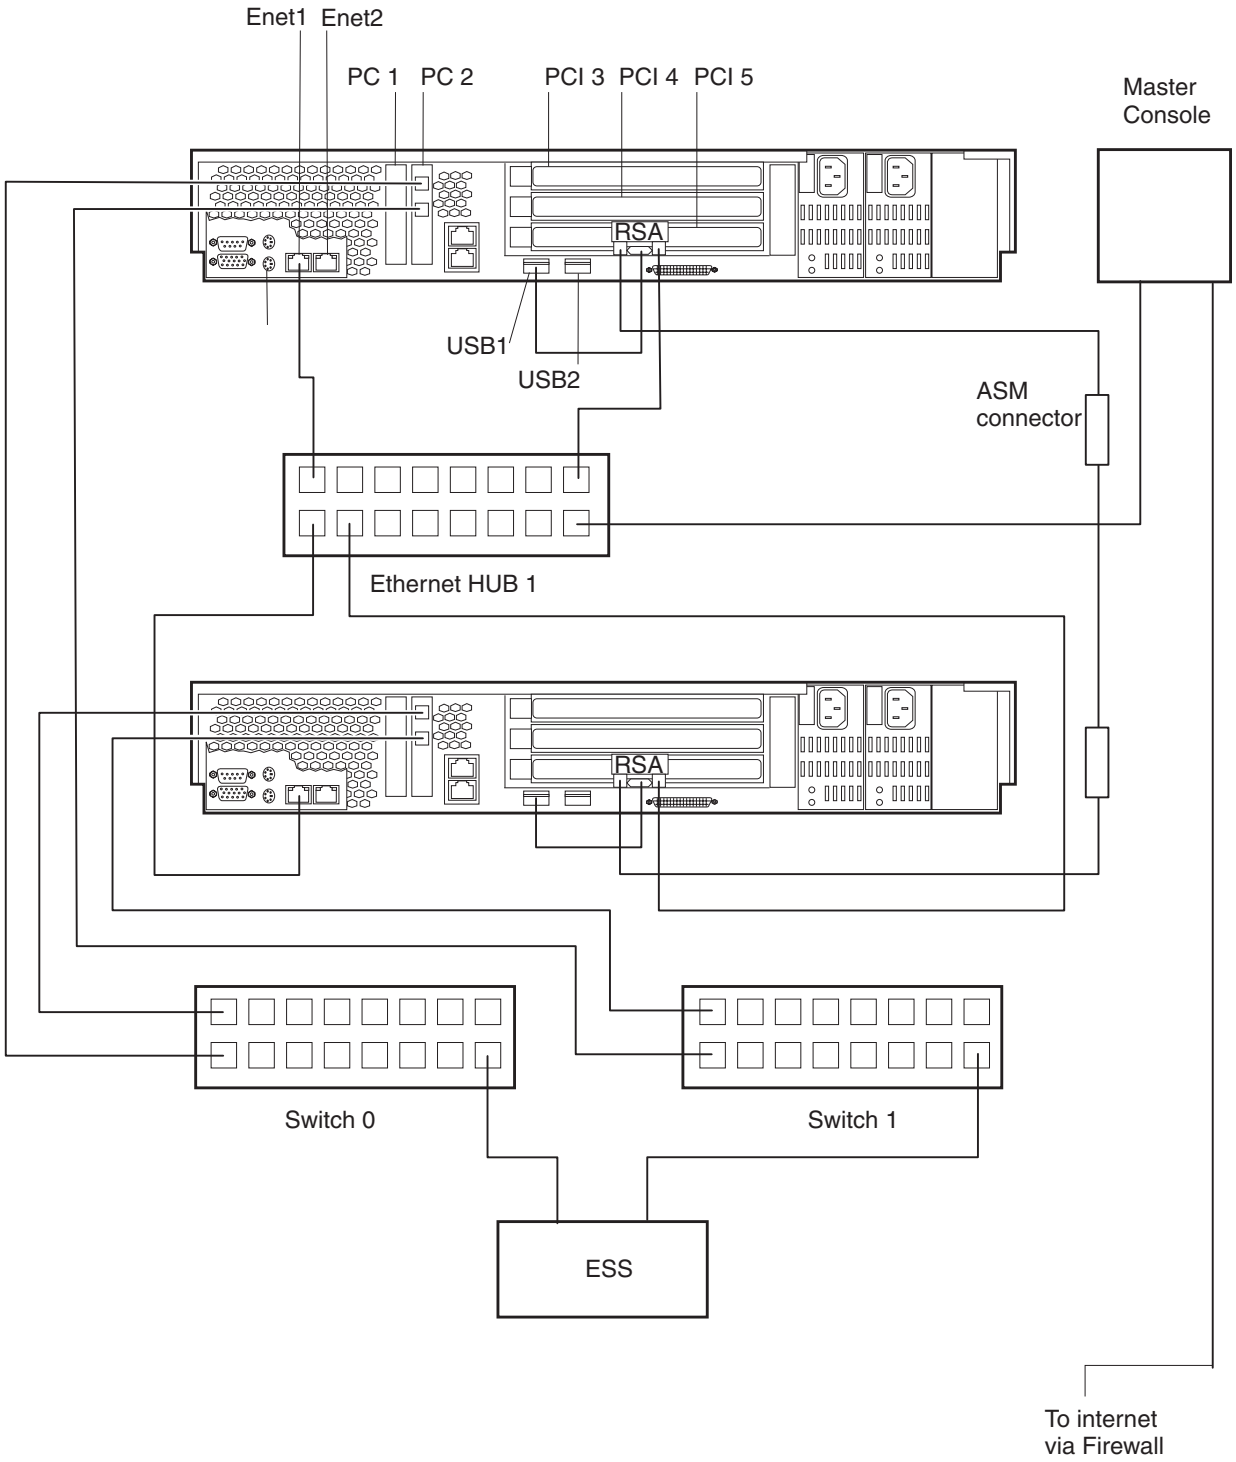

Figure 6. Two-node, two-switch, two-hub cabling example

6. Use cable clamps to secure the cables across the rear of the engine.
7. 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.

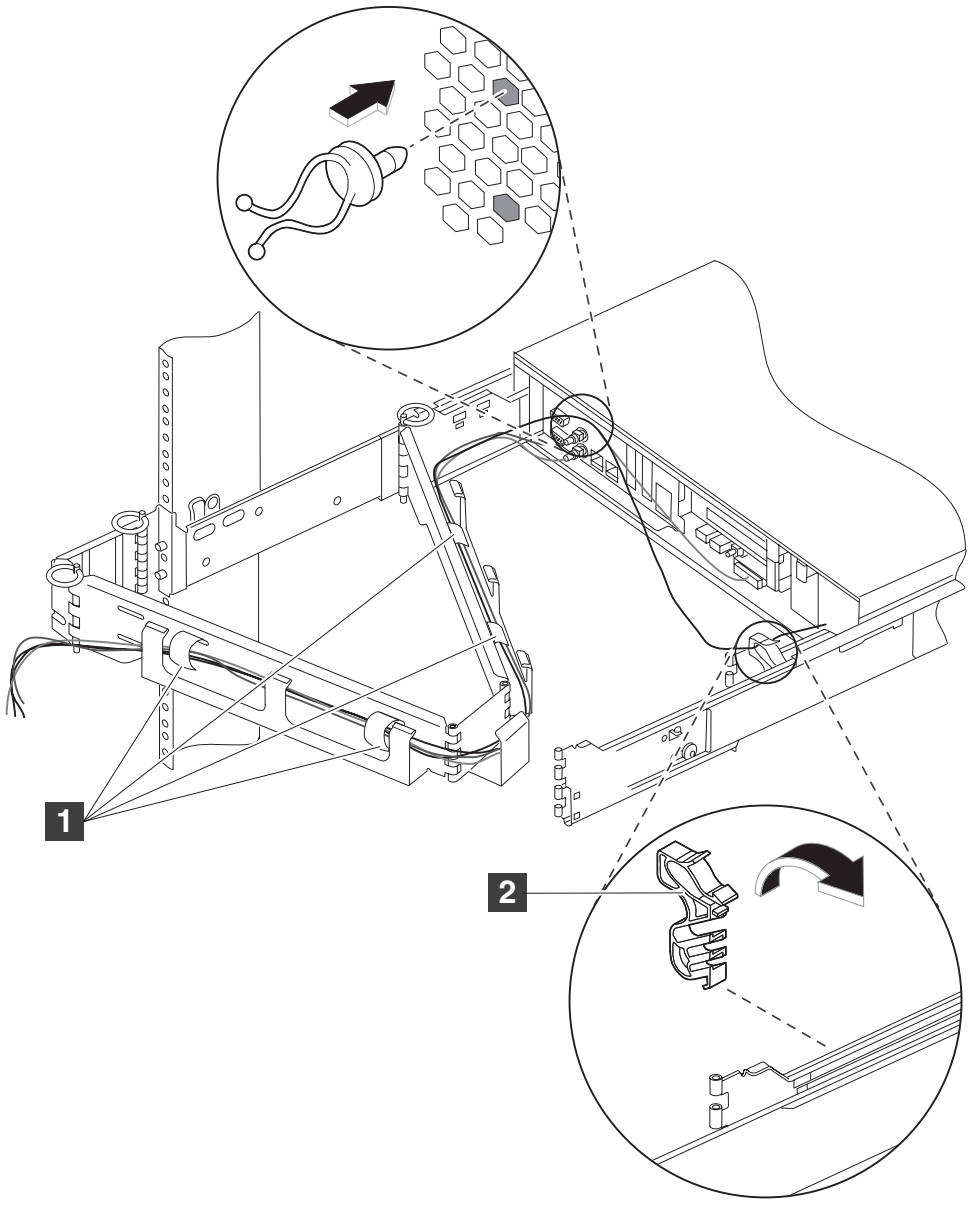

Figure 7. Attaching the cable straps

1 Cable straps
2 Cable-restraint bracket

8. 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.

For the master console, perform the following steps to cable the hardware:

1. Connect one Ethernet adapter port to the internet by way of the corporate firewall.
2. 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 6 on page 40 shows an Ethernet hub.
3. Attach the keyboard, display, and mouse to the KVM connectors on the master console.

Related topics:

- "Setting up zones" on page 29

Chapter 4. Master console

This section provides you with the necessary information for configuring the master console on a private IP network to access the SAN File System Metadata servers.

The *master console* is a serviceability node for SAN File System and other IBM TotalStorage products. SAN File System includes key features that facilitate integration with the master console. In particular, the master console provides key infrastructure for the remote access and call home features.

From the master console, you can access the following components:

- SAN File System console through a Web browser at https://<*Metadata server name*>:<*port number*> /tank where *Metadata server name* is the name or IP address of the Metadata server. The default port number is 7979.
- Administrative command-line interface through a Secure Shell (SSH) session.
- Any of the engines in the SAN File System cluster through an SSH session.
- The RSA II card for any of the engines in the SAN File System cluster through a Web browser. In addition, you can use the RSA II Web interface to establish a remote console to the engine, allowing you to view the engine desktop from the master console.
- Any of the SAN File System clients through an SSH session, a telnet session, or a remote display emulation package (such as VNC), depending on the configuration of the client.

Typically, you will use the master console to access the engines as well as the SAN File System console or the Administrative command-line interface. However, if necessary, you can attach a keyboard, monitor, and mouse to an engine.

Using the Remote Access feature of the SAN File System, you can initiate a VPN connection to allow a support engineer to remotely access the master console. You can monitor that access and disconnect the session as desired.

Related topics:

- "Software included in the SAN File System master console"
- "Accessing an engine and the Administrative command-line interface through SSH" on page 44
- "Setting up Tivoli SAN Manager (TSanM)" on page 45
- "Configuring Service Alert on the master console" on page 48
- "Setting up remote access" on page 49

Software included in the SAN File System master console

Software shipped in the master console::

- Microsoft Windows 2000 Server Edition
- IBM Director Server
- IBM Tivoli Bonus Pack for SAN Management
- FAStT Storage Manager
- PuTTY telnet/SSH package

- Adobe Acrobat Reader
- HBA Driver for the QLogic 2342 fibre-channel adapter
- VPN Connection Manager
- SAN Volume Controller Console (ICAT)

In addition to the preinstalled software listed here, the Java 2 Runtime Environment, Version 1.4 or later is required. You must download this version of the plugin as part of setting up the master console.

Accessing an engine and the Administrative command-line interface through SSH

From the master console, you can remotely access an engine using SSH. You can establish an SSH session using a password or key-pair.

Steps:

Perform these steps to establish an SSH session to any of the engines in the SAN File System cluster using a password.

1. From the master console, use one of these methods to access the engine:
   - Open a shell prompt and type **putty.exe -ssh** *engine_IP_address* where *engine_IP_address* is the IP address of the engine to be accessed.

     Note: If you used SSH to establish a remote session with the master console, type this command from that session to establish an SSH session between the master console and the engine.
   - Double-click the PuTTY icon.
     a. Fill in the IP address of the engine to be accessed.
     b. Select SSH as the protocol.
     c. Click **Open**.
2. After the session is established, log in using a Linux user ID and password.

Perform these steps to establish an SSH session to any of the engines in the SAN File System cluster using a key pair.

1. Generate SSH keys on the SSH client system by typing **ssh-keygen -b 1024 -t dsa -f** *my_key* where *my_key* is the file that will contain the private key.
2. Copy the contents of my_key.pub into a file called .ssh/authorized_keys2 in your home directory.
3. Store the SSH client public key on the SAN File System cluster.

   Note: You can use the same public key on many server machines, and your private key on many client machines. You can put several keys in your .ssh/authorized_keys2 file on the server or .ssh/identity file on the client. You can even use the same key for multiple users; for example, you might generate a system key that several users authorize to run a backup command.
4. Configure PuTTY to use keys for establishing SSH sessions:
   a. Start the PuTTY program.
   b. Click **Session** in the Category frame on the left side of the GUI.
   c. If the **SSH** radio button under the Protocol selection list on the right side of the GUI is not selected, select it.

d. Click the plus sign (+) to the left of the Connection text in the Category frame on the left side of the GUI to expand the options in this category.
e. Click SSH under the Connection category on the left side of the GUI.
f. If the radio button labeled **2** under Preferred SSH protocol version on the right side of the GUI is not selected, select it.
g. Click the plus sign (+) to the left of the SSH text in the Category frame on the left side of the GUI to expand the options in this category.
h. Click **Auth** under the SSH category on the left side of the GUI.
i. In the text box under **Private key file for authentication** on the right side of the GUI, enter the name of the file you specified to hold the client private key generated by the PuTTY key generation program. You can use the **Browse** button to the right of the field to locate the file containing the client private key. You must enter the full pathname to the file (for example, C:\Support Utils\PuTTY\priv.ppk).
j. Click **Session** in the Category frame on the left side of the GUI.
k. Click **Default Settings** in the Saved Sessions frame on the right side of the GUI.
l. Click the **Save** button to the right of the Saved Sessions selection on the right side of the GUI.

Result:

After connecting to the engine, you can access the SAN File System administrative command-line interface (CLI) to run SAN File System commands. These commands provide the ability to manage engines, Metadata servers, and Administrative servers.

Related topics:

• Chapter 4, “Master console”, on page 43

Setting up Tivoli SAN Manager (TSanM)

Before you start the TSanM, perform the following prerequisite steps to configure it:

Prerequisites:

When the master console was installed, its name might have been changed from the default name given in manufacturing to a name of your choice, if so then it will be necessary to modify some Tivoli SAN Manager configuration files to reflect the new name. Using Windows Notepad, open the following files and make the changes indicated:

1. Open the file called, c:\tivoli\itsanm\manager\bin\w32-ix86\setenv.bat. Find the line in the file *TSNM_LOCAL_HOSTNAME=xxxxxxx*, and replace the *xxxxx* with the full DNS name of the master console; then, save and close the file.
2. Open the file called, c:\tivoli\itsanm\manager\bin\w32-ix86\setDeployEnv.bat. Find the line in the file *NODE=xxxxxxx*, and replace the *xxxxx* with the Short Name for the master console; then, save and close the file.
3. Open the file called, c:\tivoli\itsanm\manager\conf\tsnmdbparms.properties. Find the line in the file *tivoli.sanmgmt.jdbc.dbURL=xxxxxxx* , and replace the *xxxxx* with the full DNS name of the master console; then, save and close the file.

4. Open the file called, c:\tivoli\itsanm\manager\conf\user.properties. Find the line in the file *SANDomainID=xxxxxxx*, and replace the *xxxxx* with the full DNS name for the master console; then, save and close the file.
5. Open the file called, c:\tivoli\itsanm\manager\apps\was\java\jre\lib\orb.properties. Find the line in the file *com.ibm.CORBA.LocalHost=xxxxxxx*, replace the *xxxxx* with the Short Name of the master console, then save and close the file.
6. Open the file called,

   c:\tivoli\itsanm\manager\apps\was\config\cells\DefaultNode\nodes\
   DefaultNode\serverindex.xml

   Find the line in the file *hostName=xxxxxxx*and replace the *xxxxx* with the master console's Short Name. Find the lines (there are 8 occurrences) *host=xxxxxxx* and replace the *xxxxx* with the master console's Short Name; then, save and close the file.
7. Open the file called, c:\WINNT\system32\drivers\etc\HOSTS. Find the last line in the file and replace the IP Address with the new address of the master console. Replace the Short Name with the new name of the master console. Replace the full DNS name with the new full DNS name of the master console, then save and close the file.

Note: To assist Tivoli SAN Manager present its optimum information, you can install the Tivoli SAN Manager Agent software on each of your host systems.

Steps:

Perform the following steps to start the TSanM:

1. Double-click **Tivoli Netview** icon on your desktop, or click **Start -> Programs –> Tivoli Console**.
2. On the menu bar, click **SAN –> Agent Configuration** and add the IP addresses of your fibre-channel switches into the SNMP Agents list.
   a. In the SNMP Agent section of the Agent Configuration panel, select each entry in turn and click **Advanced**.
   b. Type the user ID and password for that particular switch (manufacturing defaults are user ID is `Admin` and password is `passw0rd`); this is to allow TSanM to access the switch to collect zoning information.
   c. To enable this access, configure each fibre-channel switch to allow SNMP Command Access.

      Refer to your fibre-channel switch documentation for the procedure to set up this access.

   You can limit the extent of the discovery to just the components in the SAN File System. On the menu bar, click **Options –> Discovery** by editing the seed file to include the IP addresses of the fibre channel switches, the master console and IBM TotalStorage 4146s.
3. Verify the installation by running a SAN discovery. From the menu bar, click **SAN –> Configure Manager**. This displays the Configure Manager panel. Click **Clear History –> OK**. Click **Cancel** on the Configure Manager panel.

Ensure that the TSanM discovers all expected Fibre Channel connections and devices. You can visually check that the TSanM discovers all expected connections and device by displaying the topology map for each fabric and seeing that all the expected devices are presented.

Related topics:

- Chapter 4, "Master console", on page 43

Compile the Call Home MIB on the master console

Steps:

Perform the following steps to compile the Call Home MIB. Then get the RSA MIB and compile it.

1. Get the MIB at /usr/share/snmp/mibs/IBM-TANK-MIB.txt on the Metadata servers and save it on the master console as IBM-TANK-MIB.mib. One way to do this is to use PSFTP, which is preinstalled on the master console. Go to **Start > Programs > Putty > PSFTP**.
2. Open the IBM Director console.

   Note: When you log in to IBM Director, the password for the IBM Director Login window must match the password for the IBM Director Server Services.
3. In the Tasks menu, click **Discover Systems** and then **SNMP Devices**.
4. In the Groups pane on the left side of the panel, expand the All Groups group and right-click the SNMP Devices group, then click **Compile a new MIB**.
5. When the window opens, asking you to select the location of the new MIB, click the IBM-TANK-MIB.mib file that you have saved.

   The Status Messages window presents the following:

   ```
   MIB file submitted to the server.
   Starting MIB compile...
   MIB Parsing complete
   Resolving MIB imports
   Saving MIB objects...
   MIB Compile Finished.
   ```
6. Close the Status Messages window.

The RSA II adapters can be configured to send traps as well. To do this, you, download and then compile the RSA MIB. The MIB is obtained from the IBM Support Web site at http://www.ibm.com/pc/support as part of the firmware package for the RSA II adapter. You download the firmware update from IBM as a single executable file. The executable will request that you insert a diskette. Thediskette is formatted and the update software is placed on it. Then the MIB is located on the diskette.

Use the same process for compiling the RSA MIB as for the Call Home MIB.

For more information, refer to your RSA II documentation.

Related topics:

- "Configuring Service Alert on the master console" on page 48
- "Configuring Metadata servers for SNMP traps" on page 65
- "Setting up remote access" on page 49

Configuring Service Alert on the master console

Prerequisites:

The SAN File System MIB must have been compiled on the master console and the Metadata servers must have been configured to send traps to the master console.

Steps:

Perform the following steps to configure the master console for remote access.

1. Open the IBM Director console.

   Note: Verify that the IBM Director locale is set to US english.
2. In the Tasks menu, click **Event Action Plan Builder**. The Event Action Plan Builder window opens.
3. In the File menu, click **New-Simple Event Filter**. The Simple Event Filter Builder: New window opens. In this window, perform the following steps:
   - In the Event Type tab, clear the Any checkbox. In the tree on the right of the window, click **tankGenericTrap** from the tree: SNMP.iso.org.dod.internet.private.enterprises.ibm.ibmProd. ibmStorageTankModule.ibmTankTraps.
   - In the Severity tab, check the Any checkbox.
   - In the Day/Time tab, check the Any checkbox.
   - In the Category tab, clear the Any checkbox, and check the Alert category.
   - In the Sender Name tab, check the Any checkbox.
   - In the Extended Attributes tab, clear the Any checkbox. Then type the following as a single string in the Keywords edit box:

     ```
     iso.org.dod.internet.private.enterprises.ibm.ibmProd.
     ibmStorageTankModule.ibmStorageTankObjects.tankSeverity
     ```

     The operator should be ″Equal to″. The number 4 should be in the Values edit box. Click **Add** when these values have been entered.
4. In the File menu, click **Save_As** to name this filter ″IBM StorageTank Call-home Filter″.
5. In the File menu, click **New-Event Action Plan**. The Create Event Action Plan window opens. In the edit box, enter: IBM StorageTank Call-Home Filter. Click **OK**.
6. In the Event Filters view, expand the Simple Event Filter, and click the **IBM StorageTank Call-Home Filter**.
7. In the Event Action Plans view, right-click IBM StorageTank Call-Home. Choose **Add Event Filter**.
8. In the Actions view, right-click **Send an Internet (SNMP) E-mail**. In the popup, choose **Customize**.
9. Enter the following data in the Customize Action: Send an Internet (SNMP) E-mail window:
   - For Internet E-mail address: if the equipment is located in Canada, United States, Brazil, or Mexico use: callhome0@de.ibm.com; otherwise, use: callhome1@de.ibm.com.
   - For reply to: enter the E-mail address of an administrative contact at the site.

- For SMTP E-mail server: enter the name or IP address of the SMTP mail server.
- For SMTP port: enter the port number for SMTP access on the mail server.
- For Subject of E-mail Message: enter the following text: IBM StorgeTank Call-home Notification.
- For Body of E-mail message: enter the following text as shown. Spacing must be correct.

```
#machine type=xxxxxxx
#device serial number=nnnnn
#record type=1
#component id=software
#contact name=xxxxxxxxxxxxxxxxx
#contact phone=(nnn)nnn-nnnn xnnnn
#mgmt node=&system
#date recvd=&date
#time recvd=&time
```

Notes:

a. All x and n characters are to be replaced by you.
b. The machnine type is a 4-digit device type and 3-digit model. This is static for SAN File System.
c. The device serial number is the serial number of the master engine in the cluster.
d. Record type 1 indicates that this is a problem report.
e. The contact name and phone number can be any length and format.
f. The values prefixed with & must be entered as shown. They are not variables.
g. E-mails will include the text of the Call Home message following the fields presented here.

10. Save the Customized Action as IBM StorageTank Call-Home Notification.
11. In the Actions view, click the new action: **IBM StorageTank Call-Home Notification**.
12. In the Event Action Plans view, click the **IBM StorageTank Call-Home Filter**.
13. Right-click the selected action and click **Add Action** in the popup.
14. Close the Event Action Builder window.
15. Under Event Action Plans In the Tasks view of the main IBM Director Console window, drag the IBM StorageTank Call-Home Notification to the All Systems and Devices group in the Groups view.
16. When asked if you are sure, click **Yes**.

Related topics:

- "Configuring Metadata servers for SNMP traps" on page 65
- "Setting up remote access"
- "Compile the Call Home MIB on the master console" on page 47

Setting up remote access

Steps:

The following steps describe how to set up for a support representative to connect to the master console remotely through a VPN connection.

1. Log into the master console. You can access the master console directly (using the keyboard, monitor, and mouse) or remotely though another computer on the same LAN.
2. Establish a secure VPN connection from the master console through the VPN gateway to a previously designated VPN server within the IBM intranet. You establish the connection using the IBM connection manager and obtain a connection ID. The IBM connection manager icon is located on the master console desktop.
3. Provide the connection ID to the IBM support engineer. Each time you start a VPN session, a unique connection ID is created.
4. The support engineer connects to the previously designated VPN server within IBM using either a telnet client or a secure shell (SSH) client, such as PuTTY. The support engineer uses the connection ID that you provided to access the active VPN tunnel.
5. The support engineer connects to an account on the master console over the VPN connection to access to the master console. The support engineer then establishes a second connection to the VPN server. If a remote display emulation package is available, the support engineer can use this connection to establish a remote console to the master console.

Result:

This connection provides a support engineer with the ability to log on to:

- The Administrative command-line interface. The support representative can query and control the SAN File System Metadata servers, and access metadata, log, dump, and configuration data. This requires that a SAN File System administrative user account be set up for the support representative.
- Each of the engines in the SAN File System cluster. The support representative can query and control the engines at the operating system level by initiating an SSH session with the engine. This requires that a Linux user account be set up on each of the engines in the cluster.
- SAN File System clients. The support representative can query and control clients at the operating system level by initiating either an SSH session or a telnet session with the client (if an SSH or telnet application is installed and running on the client). This requires that a operating system user account be set up on each of the clients to which the support representative will need access.
- SAN File System console and RSA II Web interface (if a remote display emulation package is installed and running).

You can monitor all activity performed by the support engineer. You can either run a remote desktop package from another machine to observe the master console desktop, view the master console SSH log file to see the results of all activity, or watch directly from the monitor on the master console. In addition, you have the option to disconnect the VPN session at any time.

Related topics:

- “Configuring Service Alert on the master console” on page 48
- “Compile the Call Home MIB on the master console” on page 47
- “Configuring Metadata servers for SNMP traps” on page 65

Chapter 5. Software setup

This section provides you with the necessary instructions for setting up and configuring software on your SAN File System.

Related topics:

- "Software included in the SAN File System master console" on page 43
- Appendix H, "Managing local drives", on page 119
- "Setting up zones" on page 29
- "Installing client software" on page 66

Configuration task list

SAN File System must be prepared for the addition of new files.

Context:

As part of the configuration process, the SAN File System cluster must be prepared for the addition of new files. This includes creating new or extending existing pools, structuring the file system with filesets, and defining or activating a policy set for SAN File System. This is accomplished using the Adminstrative CLI-based interfaces.

Notes:

1. It is recommended that you issue commands using a command window that supports scrollback or sizeable displays. If the `setupTank -setmaster` command is issued from a 24-line display with no scrollback, some of the introductory information will scroll off the screen.
2. A keyboard and monitor must be attached to the node being loaded.

Steps:

Perform the following steps to prepare SAN File System for the addition of new files:

1. System configuration
   - Set date and time
   - Set IP addresses
2. Install software
   - Install WAS
   - Install server rpm
   - Install administrative rpm
3. Configure the cluster
   - *setupTank setmaster* on the master Metadata server
   - *setupTank* on the subordinate Metadata servers
   - *setupTank newcluster* on the master Metadata server
4. Configure SAN File System
   - Add volumes
   - Create pools

- create filesets

5. Install and configure clients

Related topics:

- "Configuring SAN zones" on page 30
- "Configuring LUNs" on page 16

Setting the time and date on the Metadata servers

Steps:

Set the time on the Metadata servers using the following steps:

1. Log in as root.
2. Set the clock. Example:

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

3. Set the time zone. Example:

```
bash# rm /etc/localtime
bash# ln -s /usr/share/zoneinfo/EST5EDT /etc/localtime
```

4. Set the system time from the hardware clock. Example:

```
bash# hwclock --hctosys
```

Related topics:

Assigning IP addresses for Metadata servers

In this section IP addresses are assigned.

Steps:

Perform the following steps to assign IP addresses:

1. Obtain host names and IP addresses, as well as netmask, DNS, and gateway information. Determine if any IP aliasing needs to be done.
2. Set the hostname by editing /etc/HOSTNAME. Use a single line with the fully-qualified hostname; for example: evt2-mds2.bvnssg.net.
3. Edit /etc/hosts file and add the TCP/IP address and hostname; for example: 192.168.10.41 evt2-mds2.bvnssg.net evt-mds.
4. At the command prompt, enter: **cd /etc/sysconfig/network/**.
5. Modify the IPADDR and NETMASK values in the configuration file ifcfg–eth0. See the following table for two examples.

   Note: If the engine shipped with the optional Broadcom Ethernet adapter in a PCI slot, the ifcfg-eth2 file should be modified instead. If ifcfg-eth2 does not exist, it must be created by copying the ifcfg-eth0 file to ifcfg-eth2. When this is done, rename ifcfg-eth0 and ifcfg-eth1 so that the onboard Ethernet adapters are not configured. SAN File System supports one network adapter only.

Table 2.

for eth0: (Node 0)	for eth0: (Node 1)
DEVICE=eth0 STARTMODE=onboot BOOTPROTO=static IPADDR=192.168.10.41 NETMASK=255.255.0.0	DEVICE=eth0 STARTMODE=onboot BOOTROTO=static IPADDR=192.168.10.42 NETMASK=255.255.0.0

6. If a default route needs to be added, add a line to /etc/sysconfig/network/routes. For example:

   ```
   default 192.168.1.1
   ```

7. If DNS is being used, edit the file /etc/resolv.conf. See the following example of DNS information from an existing client.

   nameserver 192.168.254.100
   nameserver 192.168.254.101
   search company.net company.com

8. Shutdown and reboot the Metadata server.
9. Login as root.
10. Use ifconfig to verify network operation by entereng the following command at the prompt: **ifconfig**.

Related topics:

- "Gather information about the existing IP network" on page 9

Setting RSA adapter IP addresses using the master console

Note: If two instances of SAN File System are to share the same master console, then the IP addresses must be assigned in a fashion that avoids assigning the same address to two different RSA adapters.

Context:

The RSA adapters in the engines all come with the same default IP address. This address is: 192.168.70.125. The recommended IP addresses for the RSA adapters in the cluster are 192.168.70.1 through 192.168.70.*n*, where *n* is incremented by 1 for each additional RSA adapter. A cluster of eight engines would contain RSA adapters numbered from 192.168.70.1 to 192.168.70.8.

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

Steps:

After all cable connections have been made, perform the following steps *for each engine*, starting with number one.

1. Power ON the engine, leaving all the others OFF.
2. On the master console open a Web browser and open the Web page at 192.168.70.125. The userID is USERID and the password is PASSW0RD. (The 0 in PASSW0RD is a zero.)
3. Ensure that the IP address of the first RSA is set to 192.168.70.1. For the first one, leave this IP address as it is. For each of the other engines' RSA adapters,

change the IP address by adding one to the last number in the IP address of each consecutive adapter. (Remember, you can only do this on one at a time.)

4. Go to the link called Network Interfaces on the side frame of the desktop and set the following:
   - Set IP interface to Static IP.
   - Set interface to Enabled.
   - Fill in values as required for IP address, subnet mask, and gateway address.
5. Click **Restart ASM** on the left sideframe.
6. Close and then reopen the browser. Log on to the browser using the *new IP address*.
7. Click **System Settings** on the left side frame.
8. Under the page for ASM information, set the value for Name. The unique name of the RSA adapter must match the name of the Metadata server on that engine. Also, set POST Watchdog to 10 minutes, OS Watchdog to 4 minutes, Loader Watchdog to 10 minutes. Click **Save**. At the bottom of the page set the time and time zone. Then click **Save**.
9. When you have finished, power this engine OFF and continue to the next one until all RSA IP addresses are set.

When you have finished, power ON all engines and perform the following steps:

1. Using your browser, go to one of the RSA II adapters, using its new IP address.
2. Log on to the adapter.
3. Click **Continue** to accept the default timeout value.
4. Click **Remote Control** on the left side of the panel.
5. If the Java1.4 plugin is not present, you will be notified. Agree to load the plugin and follow the prompts to install the plugin.

Installing the software on the engines

SAN File System software comes preinstalled on the engine. Install the other packages on **all Metadata servers** as described below.

Prerequisites:

1. Install WebSphere 5.0 Express:
   a. Login to the Metadata server as root.
   b. If you have not rebooted since the network configuration was done, set the HOSTNAME variable as follows: export `HOSTNAME=`cat /etc/HOSTNAME``.
   c. Change the directory as follows: `cd /opt/bobcat_src`.
   d. Install WebSphere 5.0 Express: `. /bobcat_install`.
   e. Repeat these steps for each Metadata server.
2. Install the administrative package on each Metadata server:
   ```
   rpm -i /usr/tank/packages/storagetank-admin-linux-1.1.0.i386.rpm
   ```
3. Install the server package on each Metadata server.:
   ```
   rpm -i /usr/tank/packages/storagetank-server-linux-1.1.0.i386.rpm
   ```

Configuration of SANFS

Use the information in the following sections to configure your SAN File System. If you want to practice with SAN File System before doing your final configuration, you can use the information provided in Appendix D, "Basic configuration for quick start", on page 95.

setupTank brings the administrative services and Metadata server clusters online upon successful completion. It can be run later to edit the Administrative server and Metadata server cluster configurations or to set their services offline. You can set many of the cluster options with the tanktool command.

You can configure parameters or sets of configuration parameters using the SAN File System console and command-line interface (CLI).

Related topics:

- Appendix D, "Basic configuration for quick start", on page 95

Creating links with device_init.sh

Use the steps in this section to create links to raw devices.

Steps:

Use the steps in this section to create links to raw devices.

1. Type `su -` to switch to a superuser.
2. Type `cd /usr/tank/server/bin` to change to the /usr/tank/server/bin directory.
3. Use `./device_init.sh` to run a script to create links to raw devices. This script should be run when LUNs have been added or removed. After setupTank has been run, device_init.sh runs automatically at each boot.

Use the steps to obtain system device names with device_init.sh. This information is required for setupTank..

1. Log in as root.
2. Enter **datapath query device** to list all vpath devices. Example: /dev/vpath.
3. Enter `device_init.sh` to create raw devices for each vpath device. Example: /dev/rvpath.

Related topics:

- "Configuring LUNs" on page 16

Starting the Metadata servers with the setupTank utility

Steps:

The setupTank utility is used to start the SAN File System configuration process.

The following steps should be complete before running setupTank.

- Chapter 3, "Physical Installation", on page 35
- "Cabling" on page 38
- "Reconfiguring Storage Devices" on page 32

- "Setting up zones" on page 29

The LDAP server should be available, and the LDAP public certificate file should be copied to /usr/tank/admin.

Notes:

1. It is recommended that you issue commands using a command window that supports scrollback or sizable displays. If the ″setupTank -setmaster″ command is issued from a 24-line display with no scrollback, some of the introductory information will scroll off the screen.
2. Before starting this process, review the list of values in the following table to be certain that you have gathered the required information from your LDAP administrator and your SAN administrator. Also, do a datapath query to determine the metadata disks.

Use the following steps:

1. Log in as root.
2. Change directory to /usr/tank/admin/bin.
3. Enter **./setupTank –setmaster**.

   The values in the following list are prompted for. Each is preceded by a description of its purpose on the screen.

Table 3.

Value	Description
SAN File System Server name	
SAN File System Cluster name	
Server IP address	
Language	
LDAP server IP address	
LDAP user (distinguished name)	
LDAP user password	
LDAP secured communication	
LDAP base distinguished name	
LDAP member attribute	
LDAP certificate	
RSA user name	The default value for RSA Username must be entered unless it has been previously changed. The default is: USERID
RSA password	The default value for RSA Password must be entered unless it has been previously changed. The default is: PASSW0RD (0 is zero)
CLI user	
CLI password	
Truststore password	
Subordinate node list	
Metadata disks	

4. To save the configuration, press **Enter** to accept the ″Yes″ default.

5. Copy /usr/tank/admin/truststore to /usr/tank/admin on each subordinate node.
6. Copy /usr/tank/admin/config/tank.properties from the master to /usr/tank/admin/config on each subordinate node.
7. On each subordinate node, run setupTank to run through all the values again. Enter new values for Server name and Server IP address, and accept the defaults for the other values. The server is now up and running.
8. Go to usr/tank/admin/bin on the master node.
9. Enter **./setupTank –newcluster**.

Configuring storage pools

Storage pool LUNs are configured as part of the installation process.

Prerequisites:

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 SAN File System as a host.
  - The administrator of the storage device must allocate LUNs to SAN File System.
- Only one system storage pool can exist.
- The storage pool name cannot already exist.
- The block size is 4 KB.
- The partition size is 16 MB.
- The maximum number of characters in the description is 256.
- The maximum number of characters in the name is 256.
- The value for alert threshold must be valid.
- You must be an administrator or IBM support representative to perform this task.

Steps:

Perform the following steps to configure storage pools:

1. Log on to the SAN File System console.
2. In the My Work pane, click **Manage Storage**.
3. Click **Storage Pools** on the portfolio to see the current set of storage pools.
4. Click **Create a Storage Pool** to start the creation process.
5. Click **Next** on the Introduction panel.
6. Set the properties for the new storage pool by specifying the following information and clicking **Next**:
   - Name of the new storage pool.
   - Description of the storage pool.
   - Partition size.
   - Allocation size.
   - Usage threshold. The default threshold is 80%.

7. Add the volumes:
   a. Specify the Volume Name Prefix.
   b. Activate the volumes by selecting the Initially Activate Volumes check box. The check box is enabled by default.
   c. Select the volumes to be added by selecting their check boxes.
   d. Click **Next**.
8. Click **Finish** to confirm the creation of the storage pool.

Related topics:

- "Configuring SAN zones" on page 30
- "Configuring LUNs" on page 16

Configuring filesets

Fileset quotas provide a way for an administrator to specify how much space can be allocated to files within a specific fileset.

Context:

Fileset quotas provide a way for an administrator to specify how much space can be allocated to files within a specific fileset. The default value is set to allow unlimited capacity; however, you can specify an alert value for the fileset. If a quota value is specified, the default alert value is 80%. If no quota value is specified, the default alert value is 0 (no alerts). When the space allocated to files within the fileset reaches the percentage of the quota, as specified by the alert value, an SNMP alert is sent to the administrator. The administrator can also specify whether to use a hard or soft quota. If a hard quota is specified, allocations that cause a quota violation will fail and an SNMP alert will be sent. If a soft quota is specified, then the allocation will be allowed to succeed and an SNMP alert will be issued.

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

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.

Steps:

Perform the following steps to configure filesets:

1. In the My Work pane, click **Manage Filing**.
2. Click **Create a Container**.

3. In the Create a Container 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.

Placement policies

Context: 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 CLI or the GUI. You cannot modify or delete it.

As part of a planning process you should consider the following:

- fileset organization
- LUN assignment to pools
- p ool allocation strategy
- assignment of filesets to servers

They are all interrelated. For example, let's say two DB2 applications running on a 8-CPU client. The executables and particularly the data files for those applications might be heavily used by the client. This, then, causes choices in how filesets are placed within the SAN File System. After all, filesets are the unit of load management, and effective choices of the filesets for the DB2 data files in particular must be made. File placement rules can build on this because fileset is one way of choosing a file's destination. Other choices might be related to the performance, availability, and capacity of the LUNs hosting the application's storage; so pool design has to factor in both the LUN's characteristics and the placement rules for the files.

The following operations may be used with respect to policies:

Create Policy Set
This operation is used to define a set of rules to be included into a new policy set.

Delete Policy Set
This operation removes a specified policy set from the Metadata server. Currently active policy sets may not be deleted.

List Policy Sets
This operation lists the policy sets that currently exist, flagging the currently activated set.

Show Policy Set
This operation retrieves the rules associated with a specified policy set.

Activate Policy Set
This operation tells the Metadata server to make an existing policy set the one that is in effect. A policy set may not be activated if any filesets or pools referenced by that policy set do not currently exist.

File placement policy syntax

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.

Notes:

1. Every policy file must start with VERSION 1.
2. The policy file must be created in seven-bit ASCII. This includes all fileset names, storage pool names, and files names that you reference in the rules.
3. A policy is not required to contain any rules, in which case it would be equivalent to the default policy.
4. 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 */).

Syntax:

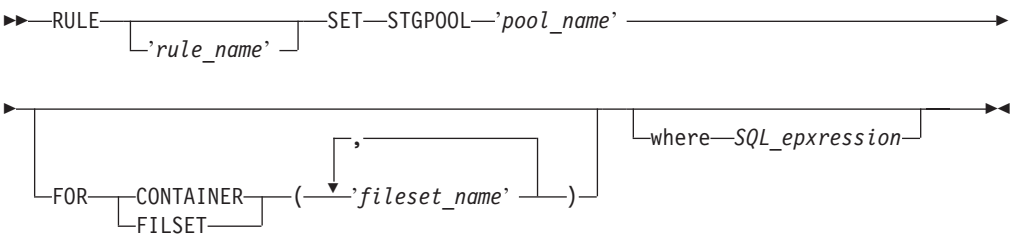

Parameters:

RULE
Initiates the rule statement.

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

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

FOR CONTAINER|FILESET ('*fileset_name***')**
Identifies one or more filesets in which the file is created to determine where the file is to be placed. In the case of nested filesets, the rules apply if the file is created in the innermost fileset.

where *SQL_expression*
Compares the file attributes specified in the rule with the attributes of the file being created to determine where the file is to be placed. The *SQL_expression* can be any combination of standard SQL-syntax expressions, including

comparison predicates, between predicates, in predicates, like predicates, mathematical value expressions, and boolean, string and numeric literals.

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

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

Attributes:

You can use any of these attributes in the expression:

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

CREATION_DATE
Date and time that the file was created.

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

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

String functions:

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

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

CHAR(*x*)
Converts an integer *x* to a string.

CHARACTER_LENGTH(*x*)
Determines the number of characters in string *x*.

CHAR_LENGTH(*x*)
Determines the number of characters in string *x*.

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 AIX-client user IDs or group IDs. These can be used in combination with comparison predicates and mathematical infix operators (such as addition, subtraction, multiplication, division, modulo division, and exponentiation).

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

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

MOD(*x*,*y*)
Determines *x* % *y*.

Date and time functions:

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

CURRENT DATE
Determines the current date on the subordinate Metadata server.

CURRENT_DATE
Determines the current date on the subordinate Metadata server

CURRENT TIME
Determines the current time on the subordinate Metadata server.

CURRENT_TIME
Determines the current time on the subordinate Metadata server.

CURRENT TIMESTAMP
Determines the current date and time on the subordinate Metadata server.

CURRENT_TIMESTAMP
Determines the current date and time on the subordinate Metadata server.

DATE(*x*)
Creates a date out of *x*.

DAY(*x*)
Creates a day of the month out of *x*.

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

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

DAYS(*x*)
Determines the number of days since 0000-00-00.

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

DAYSINYEAR(*x*)
Determines the day of the year from date *x*.

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

MINUTE(*x*)
Determines the minute from date *x*.

MONTH(*x*)
Determines the month of the year from date *x*.

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

SECOND(*x*)
Returns the seconds portion of time *x*.

TIME(*x*)
Displays *x* in a time format.

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

WEEK(*x*)
Determines the week of the year from date *x*.

YEAR(*x*)
Determines the year from date *x*.

Time and dates formats:

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

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

Timestamp
Use one of the following formats to specify a timestamp:
- *date time*
- *date*

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

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

Date Use one of these formats to specify a date:
European
DD.MM.YYYY
ISO *YYYY–MM–DD*
USA *MM/DD/YYYY*

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

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

Time Use one of these formats to specify a time:
International
HH:*MM*[*SS*[.*UUUUUU*]]
USA *HH*[:*MM*[:*SS*]] [A|P|AM|PM]

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

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

Examples:

The following example shows a sample file containing three rules that place files based on file extension.

```
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_DATE) == 1
rule 'stgRule4' set stgpool 'pool4' where USER_ID <= 100
```

Related topics:

- "Creating a policy"

Creating a policy

Prerequisites:

You must have Administrator privileges to perform this task.

Context:

SAN File System 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.

Steps:

Log onto the SAN File System console. Start the Create-policy wizard by clicking **Manage Filing ⟶ Create a Policy** in the My Work frame. Then, click **Manage Filing ⟶ Policies** in the My Work frame to verify that the policy was created.

Related topics:

- "File placement policy syntax" on page 60
- Appendix G, "Sample policy sets", on page 117

Configuring Metadata servers for SNMP traps

Prerequisites:

Use these tanktool commands on the master Metadata server to set up SNMP traps to be sent to the master console for Call Home.

Steps:

Perform the following tanktool commands on the master Metadata server:

1. `addsnmpmgr –ip` <ip address of the master console>
2. `settrap –event –sev`

Result:

Your SAN File System will now send Call Home alerts for Metadata server failures.

Related topics:

- "Configuring Service Alert on the master console" on page 48
- "Compile the Call Home MIB on the master console" on page 47
- "Setting up remote access" on page 49

Installing Subsystem Device Driver v1.4 on clients

Installation of SDD is not covered in this document. Refer to the SDD User's Guide (SC26-7540) 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.

Notes:

1. SDD is optional on client machines. If SDD is to be installed, this driver must be installed before the client package.
2. SDD comes preinstalled on SAN File System Metadata servers with the appropriate version for Linux.

Installing client software

SAN File System supports two clients: AIX v5.1 (32-bit) and Windows 2000 Server / Advanced Server.

Note: If you have some clients that use one operating system and some that use another, you must completely *install and configure* those having one operating system first; then install and configure those having the other operating system.

Steps:

Perform the following steps to install clients:

1. Ensure that the prerequisites for each type of installation are met.
2. Install the client.
3. Configure the client.

Related topics:

- "Installing the client for AIX" on page 68
- "Installing the client for Windows"

Installing the client for Windows

Prerequisites:

Note: If the system on which you are installing the SAN File System client for Windows already has another version installed, you must uninstall the existing client software before proceeding. See "Uninstalling the client for Windows" on page 87.

- The client for Windows can be installed only on Windows 2000 Server and Advanced Server. A minimum of Service Pack 4 is required. The operating system must already be installed with the appropriate service packs.
- The client for Windows requires least 10 MB of free disk space.
- You must have Administrator privileges to install the client for Windows.
- A SAN File System client can be attached to one SAN File System server cluster only, and, therefore, one global file system (GFS) 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 and manually start the client drive, the drive letter assigned to the file system will be visible only to that WTS session (private name space). To globally share the file system on WTS, the drive must be set to start automatically. This is set from a registry value under the key labelled ″STFS″, under:

  `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Stfs`
- The Metadata server must be up and running with the IP address and port defined. This information is needed during setup.
- Registry entries are built the first time SAN File System is installed. Detailed messages can be written to the internal log file pointed to by the registry entry. Severity logging is also set from the registry. Additionally, some basic startup, shutdown, or error messages are written to the Windows system log (event viewer).

Steps:

Perform the following steps to install the client for Windows:

1. If you are installing from CD, Launch the setup.exe file and follow the prompts.
2. If you are installing from the Administrative Package Repository on the Metadata server:
   a. Start the admin console on the client machine, log in, and select **Download client software** under Maintain System.
   b. Follow the instructions presented and then click **Save** in the File Download window.
   c. Going to the \temp directory and launch setup.exe.
3. Select the language to be used for the installation process, and click **OK**. The Welcome window appears.
4. Click **Next.**
5. In the Enter Configuration Parameters window, fill in the appropriate configuration information, and click **Next**.
6. Verify the information that you provided, and click **Next** to install the client driver.
7. You can have the client start automatically or manually.
   - To have the client start automatically, click **Finish** and then reboot. OR;
   - To have the client start manually:
     – Edit the Windows registry.
     – Navigate to the start key: \\HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services \STFS
     – Change the Start value to **3** by double-clicking **Start**. Then click **OK** to close the window.
     – Reboot.
8. Open a command prompt and enter **net start stfs** to start the driver.
9. Open Windows Explorer and verify that the new drive is listed.

Note: If more than one SAN File System client for Windows is running antivirus software that does directory and file scans, there is no reason to run this from more than one SAN File System client on the same files. If scans are

run from more than one client, they should be scheduled to run at different times to allow better performance of each scan.

Related topics:

- "Uninstalling the client for Windows" on page 87
- "Windows-based-client installation worksheet" on page 113
- "Tracing on the client for Windows"

Tracing on the client for Windows

Prerequisites:

Steps:

On the client for Windows, tracing is turned on by default.

1. The file size of the log is limited to 128 MB..
2. The location of the log file is: C:\program files\ibm\stfs\client\log\stfs.log
3. Only errors are enabled.

Related topics:

- "Installing the client for Windows" on page 66
- "Windows-based-client installation worksheet" on page 113

Installing the client for AIX

Prerequisites:

- You must have root privileges to install the client for AIX.
- The bos.up (for uniprocessor systems) or bos.mp (for multiprocessor systems) package must be at level 5.1.0.40 or later.
- The client for AIX must be run on a 32-bit kernel.

Context:

Notes:

1. The mount point directory must exist prior to attempting the setupstclient procedure. If it doesn't, you should "mkdir <mountpoint>" to create it.
2. All messages are written to the system log (syslog), which is set up using /etc/syslog.conf.

Steps:

Perform the following steps to install the client for AIX:

1. To determine if the AIX client software is installed enter: `lslpp –l | grep storagetank`. If you see the package name in the list, then the client is present.
2. To transfer the client package from the server repository, perform the following steps:
   a. ssh into the Metadata server.
   b. Change directory to the package repository as follows: `cd /usr/tank/packages`.
   c. ftp to the client as follows: `ftp <` .
      - bin (sets transfer mode to binary)

- cd /tmp
- put <package> (for example, put storagetank.client.aix51)
- bye (exits ftp)

3. If a file /tmp/.toc exists erase it now.
4. To install, use smit or installp to install the package. For installp, type `installp –ac –d /tmp storagetank.client.aix51` and press **Enter** to install the client.
5. To configure and start the client, go to /usr/tank/client/bin and enter `./setupstclient –prompt`.
6. Respond to the prompts that follow. (If prompted for a port number, the value entered must be consistent with the value used in the Metadata server configuration. In most cases you can accept the default.)
7. To ensure that the client is mounted, run: `ls` <mount point>.

Related topics:

- "Uninstalling the client for AIX" on page 87
- "AIX-based-client installation worksheet" on page 111

Setting administrative privileges and fileset permissions

New filesets are attached with false user ID and group ID.

Context:

New filesets are attached with false user ID and group ID for security reasons. Also ″root squashing″ makes it seem as if the user has no access.

Steps:

First, perform the following step to turn off root squashing for a client or list of clients:

1. Set the privileged client list.

```
tanktool> chclusterconfig -privclient <client_list>
```

Note: Filesets do not appear as a directory on client machines. The *attach point* is what appears, without the actual fileset name. The instructions for AIX are followed by instructions for Windows.

Steps:

For **AIX**, the following steps show you how to set permissions on a new fileset.

1. Ensure that you are logged on to a client with administrator privileges.
2. Change to the directory where the fileset is attached. This is the directory given to the ″–path″ parameter when the fileset was attached.
3. Run the ls –l command to verify that the user ID and group ID are not valid.
4. Issue the **chown** command at the client prompt to change the owner of the directory.
5. Issue the **chmod** command at the client prompt to change the permissons on the directory.

Steps:

For **Windows**, these steps show you how to set permissions on a new fileset. In this example, you set permissions on the new fileset, which was attached as ″C2″ under the root ″sanfs″.

1. Ensure that you are logged on as an administrator or a member of the administrator's group.
2. Open Windows Explorer.
3. Select the attached fileset **C2**.
4. Right-click and select **Properties**.
5. Select the **Security** tab.
6. Click **Advanced**.
7. Select the **Owner** tab.
8. The owner should be S-1-0-0, which is the null security ID.
9. Choose the proper owner, usually Administrator or Administrators.
10. Click **Apply**.
11. Click **OK**.

Verifying the installation

Steps:

To verify various parts of the installation, follow these instructions.

1. Clients: use ping to validate connectivity between between clients.
2. Metadata servers:
   a. Validating connectivity between the Metadata servers and the master console:
      - Use ping to validate connectivity.
      - Use SSH to validate connectivity.
      - Use a browser from the master console to connect to the Administrative server GUI.
   b. Show engines in the cluster: /usr/tank/admin/bin/tanktool lsserver
   c. Show clients connected to the server: /usr/tank/admin/bin/tanktool lsclient
   d. Verify that the disk devices that are intended to be used by the data and metadata storage pools are visible and accessible on each node.
   e. Validating security: you can validate the privileged clients with Isserverconfig.
   f. Validating policies: you can view the policy rules.
   g. Validating alerts and SNMP traps: you can validate the alerts set by using the Isserverconfig paramater.

Related topics:

Automatic failover

Context: The script for failover of the Metadata servers comes preinstalled on all engines, but is disabled by default. It provides failover of mastership of the cluster as well as failover of filesets in the case of irrecoverable soft and hard faults involving the loss of a single Metadata server.

The target for failover is determined by SAN File System, based on considerations such as number of nodes in the cluster and number of filesets associated with a node. The failover target is passed as an argument to the script.

The script is installed in /usr/tank/server/bin/failover.pl. Log messages are placed in /usr/tank/server/log/log.failover. If the script is to be customized, this must be done after fileset creation and assignment. The master Metadata server should not be designated as a target, and the target should not have zero or a minimum number of filesets.

Migrating data

Prerequisites:

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

Context:

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

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

Steps:

Perform the following steps to migrate your data in SAN File System:

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

Related topics:

- "Data-migration prerequisites" on page 6
- "Verifying the data integrity of migrated data"
- "Determining the data migration process" on page 21

Verifying the data integrity of 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.

Prerequisites:

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

Steps:

Perform these steps to manually verify data integrity:

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 verify** command.

Related topics:

- "Data-migration prerequisites" on page 6
- "Migrating data" on page 71

Chapter 6. Upgrading

Properties:

Note: Before performing any upgrades, you should ensure your files are backed up using your site-specific backup processes.

This section includes the following:

Upgrading the package repository

This section describes the upgrade procedure for the package repository.

The SAN File System *package repository* holds all the packages needed to install the various SAN File System software components. These include the Metadata server, the client, and the administrative server. By default, these packages are installed in /usr/tank/packages. They are bundled in package form to facilitate keeping the package repository's contents up to date. This is the primary mechanism for distributing the SAN File System software bundle.

Prerequisites:

To determine the version of the currently installed package repository, type: `rpm -qa | grep storagetank-package`.

You must have root privileges to upgrade the package repository for Linux.

Steps:

Perform the following steps to upgrade the package repository:

1. If the new package name is different from the old package name, type `rpm –e` <package_name> and press **Enter**, where *package_name* is the file name of the package repository.
2. Type `rpm –i` <package_name>, and press **Enter**, where *package_name* is the file name of the package repository. For example, if the *rpm –qa | grep storagetank-package* command returned: *storagetank-package-beta-1.0.0.i386.rpm*, then the correct command would be:

   ```
   rpm -i storagetank-package-beta-1.0.0.i386.rpm
   ```

3. Repeat for each Metadata server. Although a single up-to-date package repository can be used to serve packages to the entire SAN File System system,

for high availability purposes and to avoid accidental installation of down-level packages, it is recommended that all copies of the package repository be kept up to date.

Note: 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. To review the contents of the package before installing it, use the following comand:

```
rpm -qpl package_name
```

Related topics:

- "Upgrading the Metadata server software" on page 75
- "Upgrading clients" on page 76
- "Upgrading the Administrative server" on page 75

Metadata server software upgrades

Context:

This section covers upgrading the administrative server package and the Metadata server software.

Steps:

Perform the following steps to prepare the Metadata server for software upgrades.

1. Determine which is the master node.
   a. Change directory to: usr/tank/server/bin
   b. Issue the command: **./tanktool lscluster** to see the state of the cluster. The master console is listed as Node 0.
2. Choose a subordinate node to take offline.
   a. On the master node, go to: usr/tank/admin/bin
   b. Start tanktool: **./tanktool**
3. Issue the command: **stopserver** <server name> where *server name* is the subordinate node.
4. Enter **lsserver** to confirm the selected node is in the "Not Running" state.
5. Open an ssh session on the selected node and run the procedure to upgrade the administrative server or the procedure to upgrade the Metadata server software, or both, as required. See "Upgrading the Administrative server" on page 75 and "Upgrading the Metadata server software" on page 75.
6. To verify the upgrade, run: **rpm –qa | grep tank**
7. Restart the upgraded node. At the tanktool prompt on the master node, issue the command: **startserver <server name>** where *server name* is the name of the upgraded subordinate node.
8. Perform Step 3 through Step 7 on all remaining subordinate nodes. Then upgrade the master node by repeating the same steps for the master node.

   Note: During upgrade of the master Metadata server, the entire cluster is unavailable to the SAN.
9. When all nodes have been upgraded, commit the upgrade by issuing the following command at the tanktool prompt on the master node: **upgradecluster**

To confirm the upgrades, run **statcluster** at the master node. The entries for Software Version and Committed Software Version should match.

Related topics:

- "Upgrading the Administrative server"
- "Upgrading the Metadata server software"

Upgrading the Administrative server

An upgrade causes RPM to only replace files that have changed since they were put into place during the package installation or upgrade process. The upgrade process preserves and inherits the current configuration.

Prerequisites:

The Administrative server component should be upgraded before the Metadata server. Additionally, the entire Metadata server cluster should be upgraded before upgrading clients.

Steps:

Perform the following steps to upgrade the Administrative server:

1. Type `rpm -U` <package name>, and press **Enter** to upgrade the Administrative server component, where *package name* is the file name of the Administrative server installation package (for example, rpm -U storagetank-admin-server-RHLAS-1.0.0.rpm).

Upgrading the Metadata server software

The Metadata server component can be upgraded individually using RPM.

Note: Because installing the new operating environment overwrites SAN File System configuration files (tank.config, tank.bootstrap, and so on), be sure to back up these files to a stable external media. After the upgrade, copy back the saved files.

An upgrade causes the RPM command to only replace files that have changed since they were put into place during the package installation or upgrade process. The upgrade process preserves and inherits the current configuration.

Prerequisites:

The Administrative server component should be upgraded before the Metadata server component. Additionally, all Metadata servers in the cluster should be upgraded before upgrading clients. You must have root privileges to upgrade the Metadata server package for Linux.

Steps:

Perform the following steps to upgrade the Metadata server:

1. On the master Metadata server, type `stopserver` <server name> at the tanktool prompt.
2. Type `/usr/tank/admin/bin/stopCimom`.
3. Type `rpm –U` <package name>, where *package_name* is the file name of the Metadata server package. For example:

```
rpm -U storagetank-server-linux-1.0.0-st9_0002.i386.rpm
```

4. Type `startserver` <server name> at the prompt.
5. Type `/usr/tank/admin/bin/startCimom`.

Upgrading clients

Prerequisites: Clients should always be upgraded after the entire cluster is upgraded.

Steps:

Perform the following steps to upgrade client software:

1. Deactivate the client.
2. Uninstall the client.
3. Install the new client software.
4. Reactivate the client.

Related topics:

- "Upgrading the client for Windows"
- "Removing down-level software on a client for AIX" on page 78

Upgrading the client for Windows

Prerequisites: Clients should always be upgraded after the Metadata server cluster has been upgraded.

Document or retain a copy of the platform's registry before beginning the upgrade process.

Steps:

Begin by removing the current client software:

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. Choose **Remove**.
5. Click **Next**, then click **OK**.
6. Click **Finish**.
7. Reboot the system.

Use the following steps to install the upgraded software.

1. If you are installing from CD:
   a. Navigate to the *Windows/Win2K/Package* directory on the CD.
   b. Launch the setup.exe file.
2. If you are installing from the Administrative Package Repository, type
   a. Navigate to the */usr/tank/admin/packages/Windows/Win2K/Package* directory.
   b. Launch the setup.exe file.
3. Select the language that you want to use for the installation process, and click **OK**. The Welcome window appears.

4. Fill in the appropriate configuration information, and click **Next**.
5. Verify the information that you have provided, and click **Next** to install the client driver.
6. Click **Finish** to start the client driver.
7. Open Windows Explorer and verify that the new drive is listed.
8. Optionally, set up the driver to start automatically when the machine is booted.
9. Start SAN File System.

Related topics:

- "Installing the client for Windows" on page 66
- "Uninstalling the client for Windows" on page 87

Upgrading the client for AIX

Steps:

Perform the following steps to upgrade a client for AIX:

1. Stop all applications that use the client for AIX.
2. If the existing client is running, it is stopped and restarted automatically. However, you must first ensure that the SAN File System file system is not in use. You can do this with: **fuser –u <mount point>**. If the client is in use, the usr/tank/client/bin/stfsumount will fail.
3. Refer to the Release Notes to determine the server versions to which the client can connect. You can find the Release Notes on the Web at www.ibm.com/storage/support. Go to SAN File System > documentation.
4. To transfer the client package, perform the following steps:
   a. ssh into the Metadata server.
   b. Change directory to the package repository as follows: `cd /usr/tank/packages`.
   c. ftp to the client as follows: `ftp` <client system> .
      - `bin` (sets transfer mode to binary)
      - `cd /tmp`
      - `put` <package> (for examle, `put storagetank.client.aix51`)
      - `bye` (exits ftp)
5. Install the new client package using smit or installp (installp –ac –d <directory> <package name>). If the configuration file is invalid or incomplete (for example, a new required value was added with the upgrade) a message is displayed telling you to run setupstclient with the -prompt option. Running setupstclient with the -prompt option prompts you for all values, using the existing values as defaults. This allows you to verify that all values are correct and to enter any new values required with the upgrade.
6. To verify the upgrade, enter: `cat /usr/tank/client/VERSION`

Related topics:

- "Installing the client for AIX" on page 68
- "Uninstalling the client for AIX" on page 87

Removing down-level software on a client for AIX

Steps:

Perform the following steps to remove down-level client software:

1. Log in to the AIX client as root.
2. Ensure that the SAN File System is not is use.
3. Run **/usr/tank/client/bin/rmstclient** to unmount the client.
4. **Installp** or **smit** can be used to uninstall the client. For example: **installp –u storagetank.client.aix51**.

Related topics:

- "Installing the client for AIX" on page 68
- "Installing the client for Windows" on page 66

Chapter 7. Backing up

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 API method:

Go to "Backing up using the API method" on page 80 for information about how to perform a backup using the API-based approach.

Backing up filesets:

Go to the last step in Appendix D, "Basic configuration for quick start", on page 95 for creating FlashCopy images of selected filesets.

Related topics:

- "Backup and restore" on page 23
- "Backing up using the API method" on page 80
- "Backing up using the LUN method"
- Appendix D, "Basic configuration for quick start", on page 95

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), SAN File System clients 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.

Related topics:

- "Backup and restore" on page 23
- "Backing up using the LUN method"
- "Backing up using the API method" on page 80

Backing up using the LUN method

Prerequisites:

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 API method of backup.

Context:

Because the LUN method deals with data at the byte level, it is an all-or-nothing approach for backing up and restoring your entire SAN File System. In particular, it provides no ability to restore individual files (because it has no concept of files); you have to save and restore all the data — metadata and file data — or none of it. Restoring a previously saved FlashCopy image is the best method for recovering some subset of SAN File System data. Therefore, the LUN method is best employed as part of a disaster recovery situation.

Steps:

To back up the system using the LUN method, both the metadata and user LUNs must be in a static, consistent state. Perform the following steps to do a SAN File System backup using the LUN method:

1. Stop or pause all SAN File System client applications. Because this task is application-specific, refer to the application documentation for details on performing this step.
2. The Metadata server and all clients must complete all active transactions and flush their data to disk.

   Quiesce the SAN File System Metadata servers using the **quiescecluster -state full** command. This procedure will also lock out any subsequent new I/O from the clients or Metadata server.
3. Initiate the storage subsystem copy service using the procedure defined in its accompanying documentation.
4. After the storage subsystem copy is complete, re-enable the SAN File System Metadata servers using the **resumecluster** command.
5. Restart the client applications using the specific procedures for those applications.

A safeguard is to create backup copies of the Metadata server cluster configuration. These files: Tank.Bootstrap and Tank.Config, already have one automatic backup copy in the boot drive mirrors. This information may be regenerated from the metadata LUNs themselves, but with some difficulty. For information about creating a recovery file, refer to the *IBM TotalStorage SAN File System Administrator's Guide and Reference* on the publications CD that came with your Metadata servers.

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.

Related topics:

- "Backing up using the API method"
- "Backup and restore" on page 23
- Appendix F, "Disaster recovery", on page 115

Backing up using the API method

Prerequisites:

The API method of backup is used for SANs comprised of storage subsystems that do not offer built-in copy services. SANs that do offer copy services can use the LUN method of backup.

Steps:

You have two possible options when using the API 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 grandparent directory two levels above their original location, then follow this optimized procedure for SAN File System API backup:

During your regularly scheduled backup procedure, perform the following steps:

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 with the `mkimage` -<fileset name> -<directory name> <FlashCopy image name> command.

   This command must be executed for each fileset. Use the same <directory name> and <FlashCopy image name> for each fileset.

   Note: The term "container" is being phased out of SAN File System in favor of the term "fileset." "Container" still appears in command names, messages, and other places, although the publications use the newer term "fileset" wherever possible. The term "container" means the same as the term "fileset."
3. Save the most recent metadata to accompany the FlashCopy images with the `mkdrfile` <most recent metadata file name> command on the master Metadata server engine. Copy that file onto the client machine from which backup applications will run. Now you have a valid backup of everything you need.
4. Restart the client applications using the specific procedures for those applications.
5. Use the backup application to back up all <fileset name>/.flashcopy/<directory name> subdirectories and their contents, along with the <most recent metadata file name> file to your backup medium (usually tape).

   Attention: Backup Windows filesets only from a Windows client; backup AIX filesets only from an AIX client.

If your existing backup application does not provide the features required for the enhanced method, then follow this procedure for SAN File System API backup:

During your regularly scheduled backup procedure, perform the following steps:

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. Save the most recent metadata to accompany the FlashCopy images with the `mkdrfile` <most recent metadata file name> command on the master Metadata server engine. Copy that file onto the client machine from which backup applications will run.

3. Use the backup application to backup all <fileset name>/.flashcopy/<directory name> subdirectories and their contents, to your backup medium (usually tape).

   If possible, exclude all .flashcopy subdirectories and their contents since they will not be of any use during a subsequent restore operation.

   Note: The term "container" is being phased out of SAN File System in favor of the term "fileset." "Container" still appears in command names, messages, and other places, although the publications use the newer term "fileset" wherever possible. The term "container" means the same as the term "fileset."
4. 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.

Related topics:

- "Backing up using the LUN method" on page 79
- Appendix F, "Disaster recovery", on page 115

Saving a quick copy of your SAN File System

Steps:

To save copies of server and client information for future reference and diagnostic purposes, use the One Button Data Collector (OBDC). See "One-button data collection" on page 83 for instructions.

Related topics:

- "One-button data collection" on page 83

Saving a FlashCopy image of a fileset

Steps:

At any time you can save FlashCopy Images of filesets.

- Using the GUI, select **Maintain System** in the My Work pane, then select **Create FlashCopy Images**.
- Click **Next**.
- Under Select Containers, select the filesets for which you want a FlashCopy Image.
- Click **Next**.
- Under Set Properties, accept the defaults and then click **Next**.
- Verify your settings, then click **Finish**.
- To list the FlashCopy Images, select **Maintain System** then select **FlashCopy Images**. The default FlashCopy Image name (Image-1) should be included in the list.
- Change ownership and permissions on the fileset's .flashcopy directory to navigate it. The directory contains an entry for each FlashCopy Image name.

- Change the directory representing the image name to view the files as their images were created.

  Note: Attempting to write to the file causes an error stating that this area is read only. This applies to the .flashcopy directory and those directories and files below the .flashcopy directory.

Related topics:

- "One-button data collection"

One-button data collection

Steps:

The one-button data collection utility is designed to gather information of interest for first-failure data capture and analysis. It gathers diagnostic data of value in the initial investigation of reported problems. You should keep in mind that the amount of data gathered can be significant, so you might want to mount a file system over the output directory.

Each system provides unique log information which must be collected individually. The script must be executed on each server and each client. The script places a set of files into the directory indicated by the tool.

For servers, do the following:

1. Change to the directory where the data collector script is located as follows:
   ```
   cd /usr/tank/server/bin
   ```
2. Follow the menu options to collect the data. At the prompt enter:
   ```
   ./pmf.sh
   ```

For AIX clients, do the following:

1. Change to the directory where the data collector script is located as follows:
   ```
   cd /usr/tank/client/bin
   ```
2. Follow the menu options to collect the data. At the prompt enter:
   ```
   ./pmf.sh
   ```

For Windows clients, do the following:

1. At the C:\WINNT> prompt, change to the directory where the data collector script is located as follows:
   ```
   cd \programfiles\IBM\Storage Tank\Client\bin
   ```
2. Follow the menu options to collect the data. At the C:\WINNT> prompt enter:
   ```
   ./pmf.bat
   ```

Output is stored in C:\Program Files\Storage Tank\pmf\<system name>\Administrator

Related topics:

- "Installing the client for Windows" on page 66

Chapter 8. Uninstalling

In this chapter are the procedures for uninstalling SAN File System components.

Notes:

1. Before any uninstall, stop all running applications that use the SAN File System component that is being uninstalled.
2. If SAN File System is to be uninstalled, make sure that no applications that use SAN File System are automatically started by Windows.

Uninstalling the Package Repository:

Go to "Uninstalling the Package Repository"

Uninstalling the Metadata server:

Go to "Uninstalling the Metadata server" on page 86

Uninstalling the Administrative server:

Go to "Uninstalling the Administrative server" on page 86

Uninstalling the client for AIX:

Go to "Uninstalling the client for AIX" on page 87

Uninstalling the client for Windows:

Go to "Uninstalling the client for Windows" on page 87

Related topics:

- "Uninstalling the Administrative server" on page 86
- "Uninstalling the client for AIX" on page 87
- "Uninstalling the client for Windows" on page 87

Uninstalling the Package Repository

Configuration settings and log files are not removed when you uninstall the Package Repository.

Prerequisites:

You must have root privileges to uninstall the Package Repository for Linux.

Steps:

Perform the following steps to uninstall the Package Repository contents:

1. Type **rpm –e *installed_package*** and press **Enter** to uninstall the Package Repository contents, where *installed_package* is the file name of the Package Repository, for example:

   ```
   rpm –e storagetank-package-beta-1.0.0-i386.rpm
   ```

Note: To query the system for all installed packages with the word ″storage″ in their names, you can use the following command:

```
rpm -qa l grep storage
```

Related topics:

- “Uninstalling the Administrative server”
- “Uninstalling the client for AIX” on page 87
- “Uninstalling the Metadata server”
- “Uninstalling the client for Windows” on page 87

Uninstalling the Administrative server

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

Prerequisites:

You must have root privileges to uninstall the Administrative server for Linux.

Steps:

Perform the following steps to uninstall the Administrative server:

1. Type **rpm -e *installed_package***, and press Enter to uninstall the Administrative server component, where *installed_package* is the name of the Administrative server installation package (for example, rpm –e storagetank-admin-server-RHLAS-1.0.0.rpm).
2. To view all rpms installed on the system that have the word ″storage″ in their name, type: **rpm –ga l grep storage**.

Uninstalling the Metadata server

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

Prerequisites:

You must have root privileges to uninstall the Metadata server for Linux.

Steps:

Perform the following steps to uninstall the Metadata server:

1. Get the name of the installed package by entering: **rpm -qa | grep tank**.
2. Type **rpm –e *installed_package*** and press **Enter** to uninstall the Metadata server, where *installed_package* is the file name of the Metadata server package, for example:

Uninstalling the client for AIX

Prerequisites:

You must have root privileges to uninstall the client for AIX.

Steps:

Perform the following steps to uninstall the client for AIX:

1. To determine if the client software is running, you can enter: **/usr/bin/lpp –l |grep storagetank**.
2. The client is stopped automatically during uninstall, but if you wish to manually stop the client, type **/usr/tank/client/bin/rmstclient** and press **Enter**. This unmounts the global namespace, stops the client, and unloads the kernel module.

   Note: This utility prompts you for information necessary to remove the virtual client, and optionally saves this information in the client configuration file.
3. Type **installp –u storagetank.client.aix51** and press Enter to uninstall the client.

Related topics:

- "Installing the client for AIX" on page 68

Uninstalling the client for Windows

Prerequisites:

Steps:

Perform the following steps to uninstall the client for Windows:

1. Go to Control Panel.
2. Double-click **Add/Remove Programs**.
3. Click the IBM **SAN File System Client** in the **Currently Installed Programs** list.
4. Click **Change/Remove**.
5. In the wizard, choose **Remove** and then click **Next**.
6. Click **OK**.
7. If an option is presented that allows you to select whether or not to retain the client's existing configuration, make your choice, then click **Finish**.
8. Reboot.

Related topics:

- "Installing the client for Windows" on page 66

Appendix A. Accessibility

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 through mouse actions. 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 or Left arrows, 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 previous frame, press Shift+Ctrl+Tab.
- To print the current page or active frame, press Ctrl+P.

Appendix B. Getting help, service, and information

If you need help, service, technical assistance, or just want more information about IBM products, you will find a wide variety of sources available from IBM to assist you.

IBM maintains pages on the World Wide Web where you can get information about IBM products and services and find the latest technical information.

Table 4 lists some of these pages.

Table 4. IBM Web sites for help, services, and information

www.ibm.com/	Main IBM home page
www.ibm.com/storage/	IBM Storage home page
www.ibm.com/storage/support	IBM Support home page

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

Software Maintenance

All distributed software licenses include Software Maintenance (software subscription and technical support) for a period of 12 months from the date of acquisition providing a streamlined way to acquire IBM software and assure technical support coverage for all licenses. Extending coverage for a total of three years from date of acquisition may be elected. While your Software Maintenance is in effect, IBM will provide you assistance for your 1) routine, short duration installation and usage (how-to) questions; and 2) code-related questions. IBM provides assistance via telephone and, if available, electronic access, only to your information systems (IS) technical support personnel during the normal business hours (published prime shift hours) of your IBM support center. (This assistance is not available to your end users.) IBM provides Severity 1 assistance 24 hours a day, every day of the year.

Hardware Warranty

For a period of one year, if required, IBM provides repair or exchange service depending on the type of warranty service specified for your machine. An IBM technician will attempt to resolve your problem over the telephone; you must follow IBM's problem determination and resolution procedures. Scheduling of service will depend upon the time of your call and is subject to parts availability. Service levels are response time objectives and are not guaranteed. The specified level of warranty service may not be available in all worldwide locations; additional charges may apply outside IBM's normal service area. Contact your local IBM representative or your reseller for country and location specific information.

IBM On-Site Repair (IOR) IOR, 24 hours a day, 7 days a week, same-day response.

IBM will provide repair services for the failing machine at your location and verify its operation. You must provide suitable working area to allow disassembly and

reassembly of the IBM machine. The area must be clean, well lit, and suitable for the purpose. Depending on the proximity of the master console to the SAN File System cluster, you may also need to provide a keyboard, monitor, and mouse for attachment to a SAN File System engine.

Getting help online

Be sure to visit the support page for the SAN File System, complete with FAQs, parts information, technical hints and tips, technical publications, and downloadable files, if applicable. This page is at: www.ibm.com/storage/support.

Getting help by telephone

With the original purchase of the SAN File System, you have access to extensive support coverage. During the product warranty period, you may call the IBM Support Center (1 800 426-7378 in the U.S.) for product assistance covered under the terms of the hardware IBM warranty or the software maintenance contract that comes with product purchase.

Please have the following information ready when you call:

- Machine type and model or the SAN File System software identifier. The software identifier can be either the product name (SAN File System) or the Product Identification (PID) number.
- Serial numbers of the SAN File System engines, or your proof of purchase
- Description of the problem
- Exact wording of any error messages
- Hardware and software configuration information

If possible, have access to your computer when you call.

In the U.S. and Canada, these services are available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9:00 a.m. to 6:00 p.m. In all other countries, contact your IBM reseller or IBM marketing representative.[1]

1. Response time will vary depending on the number and complexity of incoming calls.

Appendix C. Purchasing additional services

During and after the warranty period, you can purchase additional services, such as support for other IBM and non-IBM hardware, operating systems, and application programs; network setup and configuration; extended hardware repair services; and custom installations. Service availability and name might vary by country.

Appendix D. Basic configuration for quick start

1. The first step is creating a storage pool. Here's how to do it:
   a. Using your Web browser, connect to the SAN File System console at https://<Metadata server name>:7979/tank where Metadata server name is the name or IP address of the Metadata server.
   b. In the My Work pane, click **Manage Storage –> Create a Storage Pool**.
   c. In the Create a Storage Pool wizard, click **Next**.
   d. In the Create a Storage Pool pane, fill in the Name field and the Description field under Set Properties.
   e. Optionally, select Partition Size, Allocation Size, and Usage Threshold. Click **Next**.
   f. Under Add Volumes, fill in the Volume Name Prefix field, then select a LUN from the table. Press **Next**.
   g. Verify your settings, then click **Finish**.
   h. Click **Manage Storage** –> **Storage Pools**.
   i. Verify your new storage pool in the list.
2. Creating a fileset:
   - for AIX use:
     a. In the My Work pane, click **Manage Filing** –> **Create a Container**.
     b. In the Create a Container pane, fill in the Name field and the Description field, and then select a server from the Server drop-down list.
     c. Optionally, select Quota Options and choose pertinent information.
     d. Under Attach Point, fill in the Directory Path field and the Directory Name field, then click **OK**.
     e. Verify your new fileset by selecting **Manage Filing** –> **Containers**.
     f. Grant root privileges to the client by clicking **Manage Servers and Clients** –> **Client Sessions**.
     g. In the Client Sessions pane, choose **Grant Clients Root Privileges** from the drop-down list. Click **Go**.
     h. On the IBM AIX client machine, switch to the SAN File System mount point, and change into the global fileset directory.

```
#pwd
/mnt/SAN_FS_MOUNTPT/sanfs
# ls
total 8
d---------   2 1000000  1000000    4096 Jul  3 10:21 aix51
dr-xr-xr-x   2 root    system      4096 Jul  3 10:08 .flashcopy
# |
```

     i. Change the ownership and permission of the fileset.

```
# chown root:system aix51
#chmod 755 aix51
#chown root:system aix51/.flashcopy
#chmod 555  aix51/.flashcopy
# ls
total 8
drwxr-xr-x    2 root    system   4096 Jul  3 10:21 aix51
dr-xr-xr-x    2 root    system   4096 Jul  3 10:08 .flashcopy
# |
```

     The fileset is ready for use on the AIX operating system.

- for Windows use:
  a. In the My Work pane of the GUI, click **Manage Filing** –> **Create a Container**.
  b. In the Create a Container pane, fill in the Name field and the Description field, and then select a server from the Server drop-down list.
  c. Optionally, select **Quota Options** and choose pertinent information.
  d. Under Attach Point, fill in the Directory Path field and the Directory Name field, then press **OK**.
  e. Verify your new fileset by clicking **Manage Filing** –> **Containers**. View your new fileset in the list.
  f. Grant root privileges to the client by clickting **Manage Servers and Clients** –> **Client Sessions**.
  g. In the Client Sessions pane, choose **Grant Clients Root Privileges** from the drop-down list. Click **Go**.
  h. Open Microsoft Windows Explorer and expand the SAN File System drive letter. Select the fileset you just created.
  i. Set the owner by clicking **File** –> **Properties**. Select the **Security** tab, click **Advanced**, and then click the **Owner** tab. Select an owner, click **Apply**, and then click **OK**.
  j. Set permissions by selecting the file containing the fileset. Click **File** –> **Properties**. Click the **Security** tab. Click **Advanced**, and then click the **Permissions** tab. Select a permission, then press **Apply**–> **OK** , and then press **OK** again.

3. Using a fileset for file sharing
   a. To share files on a fileset that is configured for AIX, add permissions for ″Other″ to the directory on the AIX client that represents that fileset. Those permissions map to Windows permissons for ″Everyone″. On Windows change ″Everyone″ permissions. This is reflected on AIX.
   b. After setting up a fileset for file sharing, copy some files into that fileset on one client and view them on the other client to make sure that file sharing works properly.
4. Implementing a simple policy:
   a. From the SAN File System console, click **Manage Filing** –> **Create a Policy** in the My Work pane. Click **Next**.

      In the Create a Policy pane:Under High-level settings, filll in the Name field and the Description field. Click **Next.** .
   b. Under Add Rules to Policy, fill in the **Rules Description** field with a description of the rules about .txt files (example: Move .txt files to My_New_Pool).

      Select a storage pool from the Storage Pool Assignment drop-down list (example: My_New_Pool).

      Select the **File Name** checkbox. Select **Ends with** in the drop-down list, and then fill in the adjacent field. Click **New Rule**, and repeat this step for .exe.
   c. Click **Next** when finished.
   d. In the Edit Rules for Policy pane, verify the rules. Edit if necessary, and then click **Finish**.
   e. Click **Manage Filing** –> **Policies**.

      Select the policy. From the drop-down list, select **Activate**. Click **Go**. Verify the activation.

This policy is now active. All new files created with an extension of .txt or .exe will be stored in My_New_Pool.

5. Migrating data:
   a. From the AIX client macnhine, select the directory containing the dataset to be migrated to SAN File System. This example uses /etc on an AIX client.
   b. Ensure you have root privileges by running lsclient from the ACLI.

```
tanktool> lsclient
Client    Session ID State   Server  Renewals Privilege
-----------------------------------------------------
aixclient      1     Current  ST0       83     Root
aixclient      1     Current  ST1       83     Root
```

   c. On the AIX client machine, check the space used by this set of data.

```
# du -sk /etc
8120    /etc
```

   d. Invokee the *plan* phase of this migration:

```
# /usr/tank/migration/bin/migratedata -log /tmp/log.migrate
-phase plan -destdir /mnt/SAN_FS_MOUNTPT/sanfs /etc
PLAN: Source directory: /etc
PLAN: Number of file objects to migrate: 1658
PLAN: Destination directory: /mnt/SAN_FS_MOUNTPT/sanfs/_tmp2075226185_
PLAN: On destination space required: 17.343750 MB, available: 3648 MB
PLAN: Number of CPUs: 1. Available Memory: 6 MB.  IO Blocksize: 1MB
```

   e. Invoke the *migrate* phase of this migration:

```
# /usr/tank/migration/bin/migratedata -log /tmp/log.migrate
-phase migrate -destdir /mnt/SAN_FS_MOUNTPT/sanfs /etc
PLAN: Source directory: /etc
PLAN: Number of file objects to migrate: 1658
PLAN: Destination Directory: /mnt/SAN_FS_MOUNTPT/sanfs
PLAN: On destination space required: 17.343750 MB, availoable 3648 MB
MIGRATE: Number of CPUs: 1, Available Memory: 3 MB, IO Blocksize: 1 MB
MIGRATE: COPY STARTED
MIGRATE: COPY COMPLETE: 4.703488 MB copiede at 0.145436 MB/sec
```

   f. Execute the *verify* phase of this migration:

```
# /usr/tank/migration/bin/migratedata -log /tmp/log.migrate
-phase verify -destdir
 /mnt/SAN_FS_MOUNTPD/sanfs /etc
PLAN: Source directory: /etc
PLAN: Destination directory: /mnt/SAN_FS_MOUNTPT/sanfs
VERIFY: Comparing files started.
VERIFY: SUCCEEDED: Comparing files completed with 0 errors and 0 resets
# ls
etc
# pwd
/mnt/SAN_FS_MOUNTPT/sanfs
```

6. Verifying File Placement by Policy:
   a. From the AIX client machine, create some sample .txt and .exe files. The content of these files is not examined, so they may contain anything.

```
#pwd
/mnt/SAN_FS_MOUNTPT/sanfs/aix51
#cat /etc/hosts >file1.txt
#cat /etc/hosts >file2.exe
# ls -1
total 9
-rw-r--r--      1 root     system      149 Jul  3 13:06 file1.txt
-rw-r--r--      1 root     system      149 Jul  3 13:06 file2.exe
dr-xr-xr-x      2 root     system     4096 Jul  3 10:03 .flashcopy
#
```

b. Login to the master Metadata server to view the volumes from the Administrative Command Line Interface (ACLI). Verify the volumes in the pool..

```
mds1:-# tanhitool
tanktool> lsvol
Name           State    Pool      Size (MB) Used (MB) Used %
-----------------------------------------------------------
MASTER        Activated SYSTEM        2000       192      9
My_New_Pool-1  Activated My_New_Pool   2000        16      0
```

c. Run reportvolfiles on that volume to see the list of files that are in that pool.

```
tanktool> reportvolfiles My_New_Pool-1
AIX_Fileset:aix51/file1.txt
AIX_Fileset:aix51/file2.exe
```

This shows that the sample files are in the AIX_Fileset fileset and have user data in the pool My_New_Pool-1.

7. Creating a system metadata backup:
   a. In the My Work pane of the GUI, select **Maintain System**, then select **Disaster Recovery**.
   b. In the Disaster Recovery pane select **Create** in the Recovery Files drop-down, then press **Go**.
   c. Under Create Recovery File, create a dump file by typing a file name in the Create – create new recovery file field, and then press **OK**.
   d. From the CLI, enter **lsdrfile** to check the dump files. Then enter **buildscript** to create the CLI scripts from the dump file.

```
tanktool> lsdrfile
Name     Date and Time         Size (KB)
=======================================
My_Dump Jul 23, 2003 1:01:05 AM        1
tanktool> builddrscript "My_Dump" built.
```

   e. Exit tanktool. From the bash prompt, switch to the recovery directory, /usr/tank/server/DR.

```
mds:1:~ # cd /user/tank/server/DR
mds1:/usr/tank/server/DR # ls
My_Dump.dump  TankSysCLI.attachpoint  TankSysCLI.auto  TankSysCLI.volume
```

   f. Save all four files (three scripts and one dump file). To restore your system metadata, you can use the three CLI scripts to recreate your system metadata. For more information, refer to the *IBM TotalStorage SAN File System Maintenance and Problem Determination Guide*, found on the publications CD that came with your Metadata servers.
8. Creating a FlashCopy® Image
   - From the SAN File System console click **Maintain System** –> **Create FlashCopy Images** in the My Work pane. In the Create FlashCopy Images pane, click **Next**. Under Select Containers, select the filesets for which you want a FlashCopy Image. Click **Next**. Under Set Properties, accept the defaults and then click **Next**. Verify your settings, then click **Finish**.
   - To list the FlashCopy Images, click **Maintain System** –> **FlashCopy Images**. The default FlashCopy Image name should be included in the list.
   - Change ownership and permissions on the fileset's .flashcopy directory to navigate it. The directory contains an entry for each FlashCopy Image name.

```
# pwd
/mnt/SAN_FS_MOUNTPT/sanfs/aix51
# ls -1
total 9
-rw-r--r--     1 root     system        149 Jul  3 13:06 file1.txt
```

```
-rw-r--r--     1 root     system        149 Jul  3 13:06 file2.exe
dr-xr-xr-x     3 root     system       4096 Jul  3 14:09 .flashcopy
#chown root:system  .flashcopy
#chmod 755 .flashcopy
#cd .flashcopy
# ls -1
total 1
drwxr-xr-x     2 root     system          4 Jul  3 14:09 Image-1
```

- Change the directory representing the image name to view the files as their FlashCopy images were created.

```
# cd Image-1
#ls -1
total 8
-rw-r--r--     1 root      system       149 Jul 3 13:06 file1.txt
-rw-r--r--     1 root      system       149 Jul 3 13:06 file2.exe
```

Note: Attempting to write to the file causes an error stating that this area is read only. This applies to the .flashcopy directory only.

Related topics:

- "Configuration of SANFS" on page 55

Appendix E. Worksheets

The worksheets provide you with a way to collect information to plan, configure and install SAN File System.

Worksheets:

The worksheets provided allow you to collect information for:

- SAN File System Planning
- LDAP planning
- Configuration
- Client installations

Related topics:

- "SAN File System planning worksheet" on page 102
- "Configuration worksheet" on page 104
- "LDAP planning worksheet" on page 108
- "Client installation worksheets" on page 110
- "AIX-based-client installation worksheet" on page 111
- "Windows-based-client installation worksheet" on page 113

SAN File System planning worksheet

This topic provides a worksheet for collecting information necessary to plan for the implementation of SAN File System.

Worksheet:

Use the following worksheet for collecting planning information necessary to install SAN File System.

Table 5.

Does the customer's SAN have enough switch ports and GBICs for the SAN File System servers?	_____ Yes	_____ No
Is the disk storage subsystem supported by SAN File System?	_____ Yes	_____ No
Does the customer's SAN have enough ports for the storage subsystems?	_____ Yes	_____ No
Will SAN File System be used with an existing, operational SAN?	_____ Yes	_____ No
Does the customer need SAN configuration tips? (For instance, is the customer able to zone the SAN switches properly as per SAN File System requirements? Can the customer build two separate homogeneous fabrics?)	_____ Yes	_____ No
Will SAN File System be used with an existing Ethernet LAN?	_____ Yes	_____ No
Does the customer's LAN have enough switch ports for the SAN File System servers?	_____ Yes	_____ No
Are the SAN File System clients connected to the SAN?	_____ Yes	_____ No
Are the SAN File System clients connected to the LAN?	_____ Yes	_____ No
Will the Windows® client applications work with IFS?	_____ Yes	_____ No
Will the UNIX® client applications work with VFS?	_____ Yes	_____ No
Will Windows clients access file system objects created by UNIX?	_____ Yes	_____ No
Will UNIX clients access file system objects created by Windows?	_____ Yes	_____ No
Do the customer's client systems have conflicting software?	_____ Yes	_____ No
Are any of the proposed clients also running AIX® LVM, Veritas Volume Manager, or Veritas File System?	_____ Yes	_____ No
Is the customer planning to put commercial database data under control of SAN File System?	_____ Yes	_____ No
Is the customer planning to put commercial e-mail databases under control of SAN File System?	_____ Yes	_____ No
Is Enterprise Storage Server® Subsystem Device Driver (SDD) or other path optimization software in use on the clients?	_____ Yes	_____ No
Is the operation of SAN File System to be managed with a network monitoring tool such as Tivoli® NetView®?	_____ Yes	_____ No
Does a backup/restore tool for the file data already exist?	_____ Yes	_____ No
Are LUN-oriented backup/restore facilities to be used?	_____ Yes	_____ No
Are SAN File System metadata protection facilities to be used?	_____ Yes	_____ No

Table 5. (continued)

Are there any specific performance expectations of SAN File System?	_____ Yes	_____ No
What are the resource requirements of VFS or IFS on a SAN File System client?		
How will it be known when engines should be added to the SAN File System cluster?		

Related topics:

- "Windows-based-client installation worksheet" on page 113
- "Configuration worksheet" on page 104
- "LDAP planning worksheet" on page 108

Configuration worksheet

This topic provides a worksheet for collecting information necessary to plan the configuration for SAN File System.

Worksheet:

Use the following worksheet for collecting configuration information necessary to install SAN File System.

Table 6. Worksheet

Does the customer have rack, power, and cooling for SAN File System servers? For power, 2 outlets per engine.	_____ Yes	_____ No
Is there a keyboard-video-mouse available for the server engines?	_____ Yes	_____ No
Are the storage subsystems to be managed by SAN File System of a single type?	_____ Yes	_____ No
Does the customer have an LDAP server?	_____ Yes	_____ No
Are the prospective UNIX client operating system levels supported within SAN File System?	_____ Yes	_____ No
Does the customer have a functional Network Information Service (NIS) environment?	_____ Yes	_____ No
Are the prospective Windows client operating system levels supported within SAN File System?	_____ Yes	_____ No
Is the customer using Microsoft clustering software package (MSCS)?	_____ Yes	_____ No
Does the prospective environment have a Microsoft Windows Active Directory implemented?	_____ Yes	_____ No
Does the customer have SAN File System-supported platforms, HBAs, HBA firmware, HBA drivers, switches, switch firmware?	_____ Yes	_____ No
Does the customer have switches to build two separate homogeneous fabrics (1 switch per fabric would suffice)?	_____ Yes	_____ No

Design document:

Use the following worksheet for collecting configuration information necessary to configure SAN File System.

Following the table are related suggestions for planning.

Table 7. Design document

Servers:			
Hostname	**IP 1**	**IP 2**	**RSA IP**

Amount of Required User Data Storage: ____ MB/GB/TB (circle one)

Amount of Required Metadata server Storage: ____ MB/GB/TB (circle one)

Storage Pool Summary:

Name	**Description**
SYSTEM	Metadata storage pool

Storage Summary:

Volume Name	**Size**	**Target Pool**	**Location**
		SYSTEM	

Fileset Summary:

Fileset Name	**Attach Point**	**Contents**
global fileset	/	all filesets

Policy Summary:

Set Name	**Rule Name**	**Predicate**	**Destination Pool**

Quantity of Data toMigrate: ____ MB/GB/TB (circle one)

Migration Plan Summary:

Source	**Destination**	**Migration Phase/Order**	**Approximate Size**	**Application**

Important factors to keep in mind::

- There must be enough filesets to keep all engines busy. There should be enough filesets to allow loads to be redistributed somewhat evenly after an engine failure.
- Consider whether to enforce quotas for for filesets and pools.
- Consider the performance and availablity of the LUNs that are being made available for user data and metadata.
- Consider the size of files in the system. If there is a pattern, consider controlling partition size for a pool.
- Think about the policy rules that will be introduced to take advantage of how pools are configured.
- Consider different policy sets, such as for when very different workloads occur at different times.
- Consider a strategy for dealing with unplanned spillover, such as a pool of LUNs left unassigned, or using suspended volumes.
- Consider which filesets will contain FlashCopy images, and the retention plans for each.
- Consider file sharing requirements.

MDC Audit:

Use the following worksheet for collecting configuration information necessary to configure SAN File System.

Table 8. MDC audit

Cluster Name:	
Engine Type:	
Engine OS Type/level	
MDC 1	**MDC 2**
Host name:	Host name:
BIOS level:	BIOS level:
HBA Model/driver:	HBA Model/driver:
WWPN-1:	WWPN-1:
WWPN-2:	WWPN-2:
Engine serial number:	Engine serial number:
NIC Model/driver:	NIC Model/driver:
Notes:	

Related topics:

- "Windows-based-client installation worksheet" on page 113
- "SAN File System planning worksheet" on page 102
- "LDAP planning worksheet" on page 108

LDAP planning worksheet

This topic provides a worksheet for collecting information necessary to plan for the implementation of SAN File System.

Worksheet:

Use the following worksheet for collecting planning information needed for the LDAP server.

An LDAP server is required for SAN File System. Some configuration of your LDAP server is required for SAN File System to use LDAP for authenticating SAN File System administrators. SAN File System requires an authorized LDAP username that can browse the LDAP tree where the users and roles are stored. If a secure LDAP connection is required, then the SSL certificate is needed. Fill in the tables below with your values.

Table 9. IP address, port number, username, password

Description	Recommended value	Your Value
IP address	n/a	
Port numbers	389 insecure, 636 secure	
Authorized LDAP username	n/a	
Authorized LDAP password	n/a	

Users:

A SAN File System *administrator* is the same as a *user* in the LDAP database entries. A user can use the tanktool CLI or the GUI.

Each user must have an entry in the LDAP database. All must have the same parent DN, and all must be the same objectClass. They must contain a ″user ID″ type of attribute.

Table 10. Users

Description	Recommended value	Your Value
User parent DN	ou-SANFS Users... objectclass: organizationalUnit	
objectClass of User entries	inetOrgPerson	
Attribute containing login userid	uid	

Roles:

SAN File System administrators must have an assigned role. The role determines the scope of commands which an administrator can execute. In increasing role of permissions, the four roles are: Monitor, Operator, Backup, and Administrator.

Each of the four roles must have an entry in the LDAP database. All must have the parent DN, and all must have the same objectClass. 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 role occupant's DN.

Table 11. Roles

Description	Recommended value	Your Value
Role parent DN	ou-SANFS Roles... objectclass: organizationalUnit	
objectClass of Role entries	organizationalRole	
Attribute containing role name	cn	
Attribute for role occupants	roleOccupant	

You will need to know the name of your LDAP certificate. Get this information from your LDAP administrator. The LDAP certificate is used in creating the Truststore.

Related topics:

- "Windows-based-client installation worksheet" on page 113
- "Configuration worksheet" on page 104

Client installation worksheets

The worksheets provide you with a way to collect information necessary to install SAN File System clients.

Worksheets:

The worksheets provided allow you to collect information for:

- AIX-based clients
- Windows-based clients

Related topics:

- "AIX-based-client installation worksheet" on page 111
- "Windows-based-client installation worksheet" on page 113

AIX-based-client installation worksheet

Use the following worksheet for collecting information necessary to install an AIX-based client.

Worksheet:

Client name	
Client IP address	
STFS kernel module	
Metadata server connection host	
Metadata server port	
Transport protocol	
Device-candidates list directory	
Mount file system read-only	
Display verbose messages	

Legend:

Client name

Client IP address

STFS kernel module
The client loads the file-system driver as a kernel extension. Specify the path to the location of the client kernel module file. The default is: /usr/tank/client/bin/stfs.o (This is /base/client/bin/stfs.o , where base is the base directory.)

Metadata server connection host
The fully-qualified host name or IP address of one of the Metadata servers in the cluster, in dotted decimal format (for example, 9.47.101.01).

Metadata server port number
The UDP port number of the Metadata server connection host, in dotted decimal format (for example, 10190).

Transport protocol
The transport protocol that you want the client to use to connect to the Metadata server. Specify either TCP/IP or UDP.

Device-candidates list directory
The client determines which disks to use as volumes by searching the SAN for a list of available disks, called device candidates. The device-candidate list consists of those devices that have device special files. (Device special files are UNIX files that reference hardware. The device-candidate list is not viewable to the user.) Specify the directory that contains the device special files (for example, /dev/stfsdisk/)

Mount point
The mount point (directory) from which the file system appears on the client. The default mount point is /mnt/tank.

Mount file system read-only
Specify whether you want to view, but not modify, data and metadata in the file system. The default is no.

Display verbose messages
Specify whether you want to display information messages from the commands. The default is no.

Related topics:

- "Windows-based-client installation worksheet" on page 113
- "Installing the client for AIX" on page 68

Windows-based-client installation worksheet

Use the following worksheet for collecting information necessary to install a Windows-based client.

Worksheet:

Metadata server IP address	
Metadata serverport number	
Client name	
Drive letter	

Legend:

Metadata server IP address
The IP address of one of the metadata servers in the cluster, in dotted decimal format (for example, 9.47.101.01).

Metadata server port number
The port number of the metadata server, in dotted decimal format. The default port is 1700.

Client name
The name that you want to use for the client (for example, st.ibm.com)

Drive letter
The drive letter you want to use for SAN File System storage.

Note: Windows clients should have Service Pack 3 installed.

Related topics:

- "Installing the client for Windows" on page 66
- "AIX-based-client installation worksheet" on page 111

Appendix F. Disaster recovery

Master console:

The master console comes with a recovery CD for complete restore of hard drives. Loading this CD will bring the master console back to the point where it was shipped. All configuration procedures must be repeated. Any other files must be recreated.

Metadata servers:

The disaster recovery CD for the Metadata servers reinstalls the operating system and repopulates the package repository with pristine packages. To install the Metadata server and administration packages, use the *rpm* command. Loading this CD brings the Metadata server back to the point where it was shipped. All configuration procedures must be repeated. Any other files must be recreated.

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 G. Sample policy sets

Properties:

A policy set to distribute files based on container:

```
VERSION 1
RULE 'rule1' SET STGPOOL 'pool1' FOR CONTAINER('container1','container2')
RULE 'rule2' SET STGPOOL 'pool2' FOR CONTAINER('container3')
```

A policy set to 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'
```

A policy set to distribute files based on the day of the week:

Notes:

1. The file placement resulting from this policy set can not be restored from backups.
2. This policy set assumes placement based on 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'
```

Note: Placement based on creation time, user ID, or group is not reccomended because none of these attributes work correctly when files are being restored from backup, or being migrated. In such cases, the creation time is always the time of the restore or migration, and the user and group are always those of the restore or migration application.

Related topics:

- "Creating a policy" on page 65

Appendix H. Managing local drives

Context:

It is a good idea to have SAN File System ignore local drives by including a comment line in /etc/fstab. By doing this, you can avoid configuring selected devices as part of the storage pool.

Note: Each device being masked in /etc/fstab must be presented on a separate line. You can not use a format like this:

```
#/dev/sde /dev/sdf /def/sdg
```

Instead, you must use a format like this:

```
/dev/sde
/dev/sdf
/dev/sdg
```

Steps:

Here is an example:

```
/dev/sda / ext2 defaults  1 1
#/dev/sdb /temp ext2 defaults  1 1
#exclude /dev/sdc from device_init.sh scan
# /dev/sde is not available for use by the SAN File System
```

Appendix I. IBM statement of limited warranty

Refer to the *IBM Statement of Limited Warranty* document that came with your hardware for detailed information about the country-specific warranty information.

Warranty and repair services

You can extend the hardware warranty service beyond the warranty period. Warranty and Repair Services offers a variety of post-warranty maintenance options. Availability of the services varies by product.

For more information about warranty extensions:

- In the U.S., call 1-800-426-4343.
- In Canada, call 1-800-465-7999.
- In all other countries, contact your IBM reseller or IBM marketing representative.

Appendix J. Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
MW9A/050
5600 Cottle Road
San Jose, CA 95193
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.

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	IBM	Tivoli
DB2	IBM logo	TotalStorage
Enterprise Storage Server	SecureWay	WebSphere
FlashCopy	StorageTank	xSeries

Java and all Java-based trademarks are registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks or registered trademarks of Microsoft Corporation.

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.

End of life statement

This box is a purchased unit. Therefore, it is the sole responsibility of the purchaser to dispose of it in accordance with local laws and regulations at the time of disposal.

The unit contains recyclable materials. The materials should be recycled where facilities are available and according to local regulations. In some areas, IBM may provide a product take-back program that ensure correct handling of the product. Contact your IBM representative for more information.

Safety information

The following sections describe the safety and environmental items you must consider before working with a Model 1RX.

Related topics:

- "Basic safety information (multilingual translations)"
- "General safety" on page 130
- "Electrical safety" on page 131
- "Grounding (earthing) requirements" on page 132
- "Handling electrostatic discharge-sensitive devices" on page 132
- "Handling static-sensitive devices" on page 133
- "Safety inspection guide" on page 133
- "System reliability considerations" on page 37

Basic safety information (multilingual translations)

DANGER

Before you begin to install this product, read the safety information in *Caution: Safety Information–Read This First*, SD21-0030. This booklet describes safe procedures for cabling and plugging in electrical equipment.

Gevarr: Voodrat u begint met de installatie van dit produkt, moet u eerst de veiligheidsinstructies lezen in de brochure *PAS OP! Veiligheidsinstructies–Lees dit eerst*, SD21-0030. Hierin wordt beschreven hoe u electrische apparatuur op een veilige manier moet bekabelen en aansluiten

Danger: Avant de procéder à l'nstallation de ce produit, lisez d'abord les consignes de sécurité dans la brochure *ATTENTION: Consignes de sécurité–A lire au préalable*, SD21-0030. Cette brochure décrit les procédures pour câbler et connecter les appareils électriques en toute sécurité.

Perigo: Antes de começar a instaler deste produto, leia as informações de segurança contidas em *Cuidado: Informações Sobre Segurança–Leia Primeiro*, SD21-0030. Esse folheto descreve procedimentos de segurança para a instalação de cabos e conexões em equipamentos elétricos.

危險：安裝本產品之前， 請先閱讀 "Caution: Safety Information–Read This First" SD21-0030 手冊中所提供的安全注意事項。 這本手冊將會說明使用電器設備的纜線及電源的安全程序。

Opasnost: Prije nego sto počnete sa instalacijom produkta, pročitajte naputak o pravilima o sigurnom rukovanju u Upozorenje: Pravila o sigurnom rukovanju - Prvo pročitaj ovo, SD21-0030. Ovaj privitak opisuje sigurnosne postupke za priključrivanje kabela i priključivanje na električno napajanje.

Upozornění: než zahájíte instalaci tohoto produktu, přečtěte si nejprve bezpečnostní informace v pokynech „Bezpečnostní informace“ č. 21-0030. Tato brožurka popisuje bezpečnostní opatření pro kabeláž a zapojení elektrického zařízení.

Fare! Før du installerer dette produkt, skal du læse sikkerhedsforskrifterne i *NB: Sikkerhedsforskrifter – Læs dette først* SD21-0030. Vejledningen beskriver den fremgangsmåde, du skal bruge ved tilslutning af kabler og udstyr.

Gevarr: Voordat u begint met het installeren van dit produkt, dient u eerst de veiligheidsrichtlijnen te lezen die zijn vermeld in de publikatie *Caution: Safety Information - Read This First*, SD21-0030. In dit boekje vindt u veilige procedures voor het aansluiten van elektrische appratuur.

VARRA: Ennen kuin aloitat tämän tuotteen asennuksen, lue julkaisussa *Varoitus: Turvaohjeet–Lue tämä ensin*, SD21-0030, olevat turvaohjeet. Tässä kirjasessa on ohjeet siitä, mitensähkölaitteet kaapeloidaan ja kytketään turvallisesti.

Danger : Avant d'installer le présent produit, consultez le livret *Attention : Informations pour la sécurité–Lisez-moi d'abord*, SD21-0030, qui décrit les procédures à respecter pour effectuer les opérations de câblage et brancher les équipments électriques en toute sécurité.

Vorsicht: Bevor mit der Installation des Produktes begonnen wird, die Sicherheitshinweise in *Achtung: Sicherheitsinformationen–Bitte zuerst lesen*. IBM Form SD21-0030. Diese Veröffentilchung beschreibt die Sicherheitsvorkehrungen für das Verkabien und Anschließen elektrischer Geräte.

Κίνδυνος: Πριν ξεκινήσετε την εγκατάσταση αυτού του προϊόντος, διαβάστε τις πληροφορίες ασφάλειας στο φυλλάδιο *Caution: Safety Information-Read this first*, SD21-0030. Στο φυλλάδιο αυτό περιγράφονται οι ασφαλείς διαδικασίες για την καλωδίωση των ηλεκτρικών συσκευών και τη σύνδεσή τους στην πρίζα.

Vigyázat: Mielôtt megkezdi a berendezés üzembe helyezését, olvassa el a *Caution: Safety Information–Read This First*, SD21-0030 könyvecskében leírt biztonsági információkat. Ez a könyv leírja, miyen biztonsági intézkedéseket kell megtenni az elektromos berendezés huzalozásakor illetve csatlakoztatásakor.

Pericolo: prima di iniziare l'installazione di questo prodotto, leggere le informazioni relatie alla sicurezza riportate nell'opuscolo *Attenzione: Informazioni di sicurezza–Prime informazioni da leggere* in cui sono descritte le procedure per il cablaggio ed il collegamento di apparecchiature elettriche.

危険： 導入作業を開始する前に、安全に関する
小冊子SD21-0030 の「最初にお読みください」
(Read This First)の項をお読みください。
この小冊子は、電気機器の安全な配線と接続の
手順について説明しています。

위험: 이 제품을 설치하기 전에 반드시
"주의: 안전 정보-시작하기 전에"
(SD21-0030) 에 있는 안전 정보를
읽으십시오.

ОПАСНОСТ
Пред да почнете да го инсталирате овој продукт, прочитајте
ја информацијата за безбедност:
"Предупредување: Информација за безбедност: Прочитајте го
прво ова", SD21-0030.
Оваа брошура опишува безбедносни процедури за каблирање
и вклучување на електрична опрема.

Fare: Før du begynner å installere dette produktet, må du lese sikkerhetsinformasjonen i *Advarsel: Sikkerhetsinformasjon – Les dette forst*, SD21-0030 som beskriver sikkerhetsrutinene for kabling og tilkobling av elektrisk utstyr.

Uwaga:
Przed rozpoczęciem instalacji produktu należy zapoznać się z instrukcją:
"Caution: Safety Information - Read This First", SD21-0030.
Zawiera ona warunki bezpieczeństwa przy podłączaniu do sieci elektrycznej
i eksploatacji.

Perigo: Antes de iniciar a instalação deste produto, leia as informações de segurança *Cuidado: Informações de Segurança–Leia Primeiro*, SD21-0030. Este documento descreve como efectuar, de um modo seguro, as ligações eléctricas dos equipamentos.

ОСТОРОЖНО: Прежде чем инсталлировать этот продукт, прочтите Инструкцию по технике безопасности в документе "Внимание: Инструкция по технике безопасности -- Прочесть в первую очередь", SD21-0030. В этой брошюре описаны безопасные способы каблирования и подключения электрического оборудования.

Nebezpečenstvo: Pred inštaláciou výrobku si prečítajte bezpečnosté predpisy v
Výstraha: Bezpeč osté predpisy - Prečítaj ako prvé, SD21 0030. V tejto brožúrke sú opísané bezpečnosté postupy pre pripojenie elektrických zariadení.

Pozor: Preden zaènete z instalacijo tega produkta preberite poglavje: 'Opozorilo: Informacije o varnem rokovanju-preberi pred uporabo," SD21-0030. To poglavje opisuje pravilne postopke za kabliranje,

Peligro: Antes de empezar a instalar este producto, lea la información de seguridad en *Atención: Información de Seguridad–Lea Esto Primero*, SD21-0030. Este documento describe los procedimientos de seguridad para cablear y enchufar equipos eléctricos.

Varning — livsfara: Innan du börjar installera den här produkten bör du läsa säkerhetsinformationen i dokumentet *Varning: Säkerhetsforeskrifter – Läs detta först*, SD21-0030. Där beskrivs hur du på ett säkert satt ansluter elektrisk utrustning.

危險：

開始安裝此產品之前，請先閱讀安全資訊。

注意：

請先閱讀 - 安全資訊 SD21-0030

此冊子說明插接電器設備之電纜線的安全程序。

General safety

Follow these rules to ensure general safety:

- Observe good housekeeping in the area of the machines during and after maintenance.
- When lifting any heavy object:
  1. Ensure that you can stand safely without slipping.
  2. Distribute the weight of the object equally between your feet.
  3. Use a slow lifting force. Never move suddenly or twist when you attempt to lift.
  4. Lift by standing or by pushing up with your leg muscles; this action removes the strain from the muscles in your back. *Do not attempt to lift any objects that weigh more than 16 kg (35 lb.) or objects that you think are too heavy for you.*
- Do not perform any action that causes hazards to the customer, or that makes the equipment unsafe.
- Before you start the machine, ensure that other service representatives and the customer's personnel are not in a hazardous position.
- Place removed covers and other parts in a safe place, away from all personnel, while you are servicing the machine.
- Keep your tool case away from walk areas so that other people will not trip over it.
- Do not wear loose clothing that can be trapped in the moving parts of a machine. Ensure that your sleeves are fastened or rolled up above your elbows. If your hair is long, fasten it.
- Insert the ends of your necktie or scarf inside clothing or fasten it with a nonconductive clip, approximately 8 centimeters (3 inches) from the end.
- Do not wear jewelry, chains, metal-frame eyeglasses, or metal fasteners for your clothing.

  Remember: Metal objects are good electrical conductors.
- Wear safety glasses when you are: hammering, drilling soldering, cutting wire, attaching springs, using solvents, or working in any other conditions that might be hazardous to your eyes.
- After service, reinstall all safety shields, guards, labels, and ground wires. Replace any safety device that is worn or defective.

- Reinstall all covers correctly before returning the machine to the customer.

Electrical safety

CAUTION:
Electrical current from power, telephone, and communication cables can be hazardous. To avoid personal injury or equipment damage, disconnect the attached power cords, telecommunication systems, networks, and modems before you open the appliance covers, unless instructed otherwise in the installation and configuration procedures.

Observe the following rules when working on electrical equipment.

Attention: Use only approved tools and test equipment. Some hand tools have handles covered with a soft material that does not insulate you when working with live electrical currents.

Many customers have, near their equipment, rubber floor mats that contain small conductive fibers to decrease electrostatic discharges. Do not use this type of mat to protect yourself from electrical shock.

- Find the room emergency power-off (EPO) switch, disconnecting switch, or electrical outlet. If an electrical accident occurs, you can then operate the switch or unplug the power cord quickly.
- Do not work alone under hazardous conditions or near equipment that has hazardous voltages.
- Disconnect all power before:
  - Performing a mechanical inspection
  - Working near power supplies
  - Removing or installing main units
- Before you start to work on the machine, unplug the power cord. If you cannot unplug it, ask the customer to power-off the wall box that supplies power to the machine and to lock the wall box in the off position.
- If you need to work on a machine that has exposed electrical circuits, observe the following precautions:
  - Ensure that another person, familiar with the power-off controls, is near you.

    Remember: Another person must be there to switch off the power, if necessary.
  - Use only one hand when working with powered-on electrical equipment; keep the other hand in your pocket or behind your back.

    Remember: There must be a complete circuit to cause electrical shock. By observing the above rule, you may prevent a current from passing through your body.
  - When using testers, set the controls correctly and use the approved probe leads and accessories for that tester.
  - Stand on suitable rubber mats (obtained locally, if necessary) to insulate you from grounds such as metal floor strips and machine frames.

Observe the special safety precautions when you work with very high voltages; these instructions are in the safety sections of maintenance information. Use extreme care when measuring high voltages.

- Regularly inspect and maintain your electrical hand tools for safe operational condition.
- Do not use worn or broken tools and testers.
- *Never assume* that power has been disconnected from a circuit. First, *check* that it has been powered off.
- Always look carefully for possible hazards in your work area. Examples of these hazards are moist floors, nongrounded power extension cables, power surges, and missing safety grounds.
- Do not touch live electrical circuits with the reflective surface of a plastic dental mirror. The surface is conductive; such touching can cause personal injury and machine damage.
- Do not service the following parts with the power on when they are removed from their normal operating places in a machine:
  - Power supply units
  - Pumps
  - Blowers and fans
  - Motor generators
  - Similar units

  This practice ensures correct grounding of the units.
- If an electrical accident occurs:
  - Use caution; do not become a victim yourself.
  - Switch off power.
  - Send another person to get medical aid.

Grounding (earthing) requirements

Electrical grounding (earthing) of the computer is required for operator safety and correct system function. Proper grounding of the electrical outlet can be verified by a certified electrician.

Handling electrostatic discharge-sensitive devices

Any computer part containing transistors or integrated circuits (ICs) should be considered sensitive to electrostatic discharge (ESD). ESD damage can occur when there is a difference in charge between objects. Protect against ESD damage by equalizing the charge so that the machine, the part, the work mat, and the person handling the part are all at the same charge.

Notes:

1. Use product-specific ESD procedures when they exceed the requirements noted here.
2. Make sure that the ESD protective devices you use have been certified (ISO 9000) as fully effective.

When handling ESD-sensitive parts:

- Keep the parts in protective packages until they are inserted into the product.
- Avoid contact with other people.
- Wear a grounded wrist strap against your skin to eliminate static on your body.
- Prevent the part from touching your clothing. Most clothing is insulative and retains a charge even when you are wearing a wrist strap.

- Use the black side of a grounded work mat to provide a static-free work surface. The mat is especially useful when handling ESD-sensitive devices.
- Select a grounding system, such as those listed below, to provide protection that meets the specific service requirement.

  Note: The use of a grounding system is desirable but not required to protect against ESD damage.

  – Attach the ESD ground clip to any frame ground, ground braid, or green-wire ground.
  – Use an ESD common ground or reference point when working on a double-insulated or battery-operated system. You can use coaxial or connector-outside shells on these systems.
  – Use the round ground-prong of the AC plug on AC-operated computers.

Handling static-sensitive devices

Attention: Static electricity can damage electronic devices and your engine. To avoid damage, keep static-sensitive devices in their static-protective packages until you are ready to install them.

To reduce the possibility of electrostatic discharge, observe the following precautions:

- Limit your movement. Movement can cause static electricity to build up around you.
- Handle the device carefully, holding it by its edges or its frame.
- Do not touch solder joints, pins, or exposed printed circuitry.
- Do not leave the device where others can handle and possibly damage the device.
- While the device is still in its static-protective package, touch it to an unpainted metal part of the server for at least 2 seconds. (This drains static electricity from the package and from your body.)
- Remove the device from its package and install it directly into the server without setting it down. If it is necessary to set the device down, place it in its static-protective package. Do not place the device on the engine cover or on a metal table.
- Take additional care when handling devices during cold weather because heating reduces indoor humidity and increases static electricity.

Related topics:

- "Handling electrostatic discharge-sensitive devices" on page 132

Safety inspection guide

The intent of this inspection guide is to assist you in identifying potentially unsafe conditions on these products. Each machine, as it was designed and built, had required safety items installed to protect users and service personnel from injury. This guide addresses only those items. However, good judgment should be used to identify potential safety hazards due to attachment of non-IBM features or options not covered by this inspection guide.

If any unsafe conditions are present, you must determine how serious the apparent hazard could be and whether you can continue without first correcting the problem.

Consider these conditions and the safety hazards they present:
- Electrical hazards, especially primary power (primary voltage on the frame can cause serious or fatal electrical shock)
- Explosive hazards, such as a damaged CRT face or bulging capacitor
- Mechanical hazards, such as loose or missing hardware

The guide consists of a series of steps presented in a checklist. Begin the checks with the power off and the power cord disconnected.

Checklist:

1. Check exterior covers for damage (loose, broken, or sharp edges).
2. Power OFF the engine. Disconnect the power cord.
3. Check the power cord for:
   a. A third-wire ground connector in good condition. Use a meter to measure third-wire ground continuity for 0.1 ohm or less between the external ground pin and frame ground.
   b. The power cord should be the appropriate type as specified in the parts listings.
   c. Insulation must not be frayed or worn.
4. Remove the cover.
5. Check for any obvious non-IBM alterations. Use good judgment as to the safety of any non-IBM alterations.
6. Check inside the unit for any obvious unsafe conditions, such as metal filings, contamination, water or other liquids, or signs of fire or smoke damage.
7. Check for worn, frayed, or pinched cables.
8. Check that the power-supply cover fasteners (screws or rivets) have not been removed or tampered with.

Related topics:
- "General safety" on page 130
- "Safety information" on page 125

Electronic emission notices

Related topics:
- "Federal Communications Commission (FCC) statement"
- "Australia and New Zealand Class A statement" on page 135
- "Industry Canada Class A emission compliance statement" on page 135
- "Chinese Class A warning statement" on page 136
- "European Union EMC Directive conformance statement" on page 135
- "Japanese Voluntary Control Council for Interference (VCCI) statement" on page 136
- "Taiwan electrical emission statement" on page 136
- "United Kingdom telecommunications safety requirement" on page 135

Federal Communications Commission (FCC) statement

Federal Communications Commission (FCC) Class A Statement:

Properly shielded and grounded cables and connectors must be used in order to meet FCC emission limits. IBM is not responsible for any radio or television interference caused by using other than recommended cables and connectors or by unauthorized changes or modifications to this equipment. Unauthorized changes or modifications could void the user's authority to operate the equipment.

This device complies with Part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) this device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operation.

Note: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case the user will be required to correct the interference at his or her own expense.

Industry Canada Class A emission compliance statement

This Class A digital apparatus complies with Canadian ICES-003.

Avis de conformité à la réglementation d'Industrie Canada

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

Australia and New Zealand Class A statement

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

United Kingdom telecommunications safety requirement

Notice to Customers

This apparatus is approved under approval number NS/G/1234/J/100003 for indirect connection to public telecommunication systems in the United Kingdom.

European Union EMC Directive conformance statement

This product is in conformity with the protection requirements of EU Council Directive 89/336/EEC on the approximation of the laws of the Member States relating to electromagnetic compatibility. IBM cannot accept responsibility for any failure to satisfy the protection requirements resulting from a nonrecommended modification of the product, including the fitting of non-IBM option cards.

This product has been tested and found to comply with the limits for Class A Information Technology Equipment according to CISPR 22/European Standard EN 55022. The limits for Class A equipment were derived for commercial and industrial environments to provide reasonable protection against interference with licensed communication equipment.

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Chinese Class A warning statement

Attention: This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

中华人民共和国“A类”警告声明

声 明

此为A级产品，在生活环境中，该产品可能会造成无线电干扰。在这种情况下，可能需要用户对其干扰采取切实可行的措施。

Taiwan electrical emission statement

警告使用者:
這是甲類的資訊產品，在居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求採取某些適當的對策。

Japanese Voluntary Control Council for Interference (VCCI) statement

この装置は、情報処理装置等電波障害自主規制協議会（ＶＣＣＩ）の基準に基づくクラスＡ情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

Glossary

This glossary includes terms and definitions from:

- *The American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies can be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *ANSI/EIA Standard - 440A: Fiber Optic Terminology*, copyright 1989 by the Electronics Industries Association (EIA). Copies can be purchased from the Electronics Industries Association, 2001 Pennsylvania Avenue N.W., Washington, D.C. 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.
- The *Storage Networking Dictionary*, available online at the Storage Networking Industry Association (SNIA) Web site:

  www.snia.org/education/dictionary/
- The Distributed Management Task Force (www.dmtf.org), copyright 2003 by the Distributed Management Task Force, Inc., 225 SE Main Street Portland, OR 97214. Definitions derived from this book have the symbol (D) after the definition.

This glossary uses the following cross-reference forms:

See This refers the reader to one of two kinds of related information:

- A term that is the expanded form of an abbreviation or acronym. This expanded form of the term contains the full definition.
- A synonym or more preferred term

See also This refers the reader to one or more related terms.

ACLI. See *Administrative command-line interface (ACLI)*.

Administrative command-line interface (ACLI). A command-line interface used to administer all aspects of the SAN File System. The ACLI runs on all engines that host Metadata servers and the Administrative server.

administrative log. A log that maintains a history of messages created by the Administrative server.

Administrative server. For SAN File System, a set of servlets running within a customized instance of WebSphere Application Server that handles all SAN File System administrative requests from the SAN File System console. See also *SAN File System console*.

alert. A message or other indication that identifies a problem or an impending problem.

audit log. A log that maintains the history of all commands issued by any administrator for all Metadata servers in the cluster.

CIM. See *Common Information Model*.

CIM client application. A storage management program that initiates CIM requests to the Administrative agent for the device.

CIM namespace. The scope within which a CIM schema applies.

CIM object manager (CIMOM). The common conceptual framework for data management that receives, validates, and authenticates the CIM requests from the client application and then directs the requests to the appropriate component or device provider.

CIMOM. See *CIM object manager*.

client. For SAN File System, a client is a system that can access the SAN File System. These clients act as servers to a broader clientele, providing Network File System or Common Internet File System access to the global namespace or hosting applications (such as database servers or Web-hosting services that use multiple servers).

class. The definition of an object within a specific hierarchy. An object class can have properties and methods and serve as the target of an association.

CLI. See *Administrative command-line interface*.

client state manager (CSM). A component of the client kernel that provides protocol support for the client.

cluster. A group of engines that is managed as a set and presents a single point of control for configuration and service activity.

cluster log. A log that maintains a history of messages created by all Metadata servers in the cluster.

cluster state. A status condition of the cluster. Cluster states can be inactive (Not running or Forming), active (Online, Offline, Partly quiescent, or Fully quiescent) or unknown. See also *Forming, Fully quiescent, Not running, Offline, Online,* and *Partly quiescent*.

Common Information Model (CIM). A set of standards from the Distributed Management Task Force Inc. (DMTF). CIM provides a conceptual framework for storage management and an open approach to the design and implementation of storage systems, applications, databases, networks, and devices.

coordinated universal time (UTC). The time scale, based on the System International (SI) second, as defined and recommended by the Comitb Consultatif International de la Radio (CCIR) and maintained (using an atomic clock) by the Bureau International des Poids et Mesures (BIPM).

CSM. See *client state manager*.

default user storage pool. A storage pool that stores file data that SAN File System has not assigned (using the active policy) to a user storage pool, as well as file data that is assigned directly to this storage pool. There is only one default user storage pool; however, you can assign any user storage pool as the default storage pool. See also *user storage pool*

engine. The hardware unit that hosts the software for the Metadata server.

event log. (1) A log that maintains a history of event messages issued by all Metadata servers in the cluster. (2) IBM Term: A log that contains information about events for a particular system or group, for a particular metric, or for all the events that are associated with a specific monitor.

file-placement rule. A rule that controls in what pool SAN File System places files in the global namespace. See also *rule* and *global namespace*.

fileset. A hierarchical grouping of files managed as a unit for balancing workload across a cluster.

FlashCopy image. A space-efficient image of the contents of part of the SAN File System at a particular moment.

Forming. A status condition where the cluster has a master and is in the process of forming. This state is always the initial one whenever a cluster is newly formed.

Fully quiescent. A status condition that cuts off all client communication with the cluster.

global namespace. A single file system that provides complete, shared access to both Windows and UNIX clients in the same environment.

ID. See *identifier*.

identifier (ID). A sequence of bits or characters that identifies a user, program, device, or system to another user, program, device, or system.

Initializing. A status condition during which a Metadata server or the entire cluster is set up for the first time.

key. A property that is used to provide a unique identifier for an instance of a class. Key properties are marked with the Key qualifier. (D)

lease. The amount of time that a client can hold a lock.

lock. A restriction that allows clients to have exclusive access to files. Types of locks include *data locks*, *session locks*, and *range locks*.

logical unit (LU). In open systems, a logical disk drive.

logical unit number (LUN). In the small computer system interface (SCSI) protocol, a unique number used on a SCSI bus to enable it to differentiate between up to sixteen separate devices per SCSI ID address, each of which is a logical unit.

LU. See *logical unit*.

LUN. See *logical unit number*.

managed object format (MOF). A compiled language for defining classes and instances. A MOF compiler offers a textual means of adding data to the CIM Object Manager repository. MOF eliminates the need to write code, thus providing a simple and fast technique for modifying the CIM Object Manager repository. (D)

master Metadata server. In SAN File System, the Metadata server in a cluster that is responsible for physical-space allocation.

metadata. Data that describes the characteristics of stored data; descriptive data.

Metadata server. In SAN File System, a server that offloads the metadata processing from the data-storage environment to improve SAN performance. An instance

of the SAN File System runs on each engine, and together the Metadata servers form a cluster. See also *cluster*.

method. A way to implement a function on a class.

MOF. See *managed object format*.

Not running. A status condition where one or more servers in the cluster are not added and therefore the cluster cannot perform any functions.

object name. An object that consists of a CIM namespace path and a model path. The namespace path provides access to the CIM implementation managed by the CIM agent, and the model path provides navigation within the implementation.

Offline. A status condition during which clients are not being serviced and the cluster is responding only to administrative requests.

Online. A status condition that indicates the normal operational state for the cluster.

Partly Quiesced. A state in which the cluster or server is in a "quiet" client communications mode to allow other operations to occur.

Partly quiescent. A status condition that allows existing metadata activity and client communication to continue on the cluster, but prohibits new communication.

policy. A list of file-placement rules that define characteristics and placement of files. Several policies can be defined within the configuration, but only one policy is active at one time. See also *file-placement rule* and *service-class rule*.

pool. See *storage pool*.

property. An attribute that is used to characterize instances of a class.

qualifier. A value that provides additional information about a class, association, indication, method, method parameter, instance, property, or reference.

quota. A limit on the amount of disk space a user can use.

rule. The lines within a policy that specify which actions will occur when certain conditions are met. Conditions include attributes about an object (file name, type or extension, dates, owner, and groups) and the fileset name associated with the object.

SAN File System console. A Web user interface used to monitor and control the SAN File System remotely by using any standard Web browser.

schema. A group of object classes defined for and applicable to a single namespace. Within the CIM agent, the supported schemas are loaded through the managed object format (MOF) compiler.

security log. A log that maintains a history of Administrative server login activity.

service location protocol. A directory service that the CIM client application calls to locate the CIMOM.

Shutdown. A status condition that describes when the cluster is shut down as intended.

SLP. See *service location protocol*.

Starting. A status condition when a Metadata server is starting as designed but is not ready to accept connections from clients.

storage pool. A named set of storage volumes that is the destination for storing client data.

symbolic link. A type of file that contains the path name of and acts as a pointer to another file or directory.

system storage pool. A storage pool that 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. See also *Metadata server*.

user storage pool. An optional storage pool that contains blocks of data that compose the files that are created by SAN File System clients. See also *storage pool* and *default user storage pool*.

volume. A labeled logical unit, which can be a physical device or a logical device. For SAN File System, there is a one to one relationship between volumes and LUNs. See also *logical unit number*.

UTC. See *coordinated universal time*

Index

M

N

O

P

R

S

T

U

V

W

Z

Readers' Comments — We'd Like to Hear from You

IBM TotalStorage™ SAN File System
(based on IBM Storage Tank™ technology)
Planning, Installation and Configuration Guide
Version 1 Release 1

Publication No. GA27-4316-00

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
GA27-4316-00

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

GA27-4316-00

Cut or Fold
Along Line

Part Number: 17P7093

Printed in U.S.A.

(1P) P/N: 17P7093

GA27-4316-00