Tivoli

Using Tivoli Storage Manager in a SAN Environment

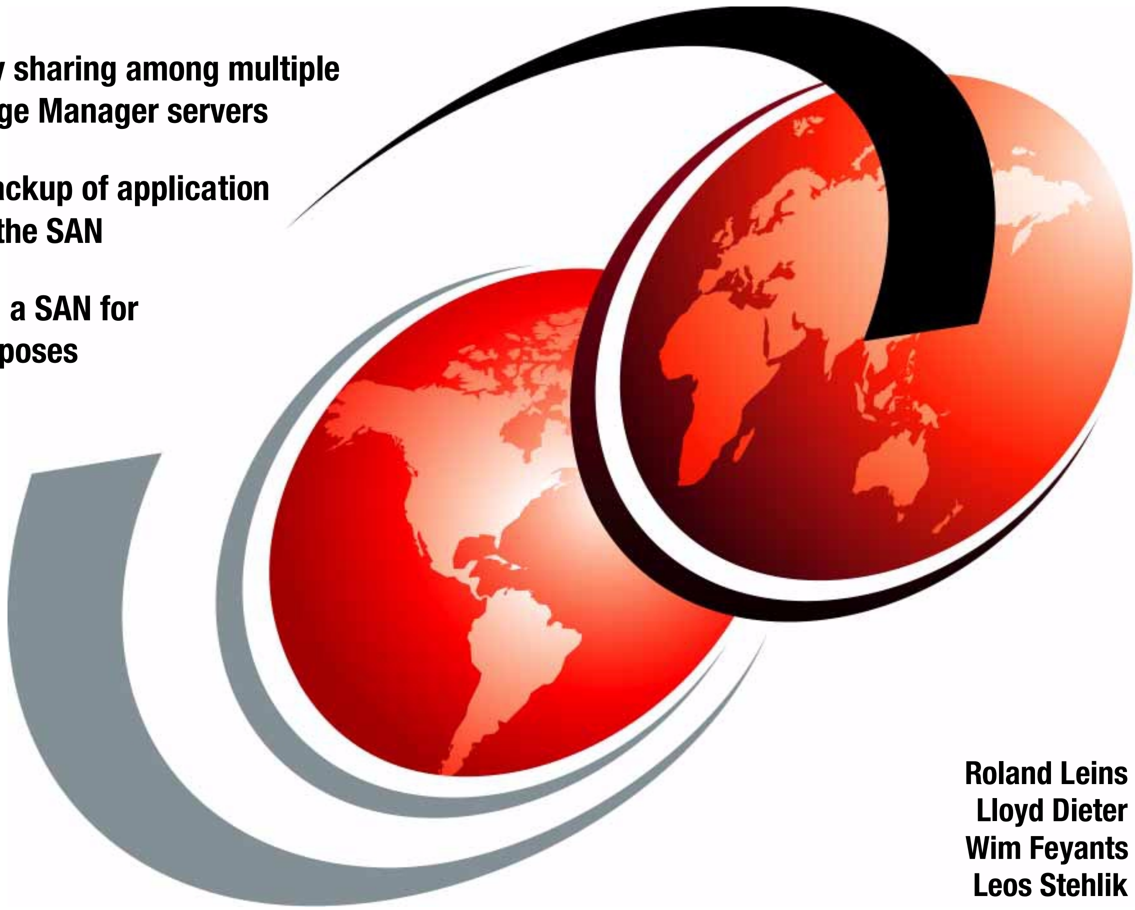

- Tape library sharing among multiple Tivoli Storage Manager servers
- LAN-free backup of application data using the SAN
- Configuring a SAN for backup purposes

Roland Leins
Lloyd Dieter
Wim Feyants
Leos Stehlik

Redbooks

ibm.com/redbooks

SG24-6132-00

International Technical Support Organization

Using Tivoli Storage Manager in a SAN Environment

December 2000

Take Note!

Before using this information and the product it supports, be sure to read the general information in Appendix C, “Special notices” on page 281.

First Edition (December 2000)

This edition applies to Version 4, Release 1 of Tivoli Storage Manager, Program Numbers 5698-TSM for use with the Microsoft Windows NT and Windows 2000, IBM AIX, Sun Solaris, HP-UX, and IBM OS/400, and to Version 2, Release 1 of Tivoli SANergy File Sharing, Program Number 5698-SFS for use with Microsoft Windows NT, Windows 2000, and Sun Solaris.

Comments may be addressed to:
IBM Corporation, International Technical Support Organization
Dept. 471F Building 80-E2
650 Harry Road
San Jose, California 95120-6099

Contents

Figures

Tables

Preface

Storage Area Networks (SANs) are changing the way data is moved, stored, and managed, and they stand to transform the computer landscape in much the same way that LANs did in the 1980s. Tivoli Storage Manager has been enhanced in the latest releases to provide new functionality that utilizes SAN technology.

This IBM Redbook will help give a broad understanding of SANs, and will provide practical implementation details for using Tivoli Storage Manager products to tap the new speed, flexibility, and manageability afforded by a SAN.

This book is written for IBM, Tivoli, customer, vendor, and consulting personnel who wish to gain an understanding of how SAN technology applies to a Tivoli Storage Manager environment, or to obtain assistance in implementing Tivoli Storage Manager utilizing SAN-attached storage.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization San Jose Center.

Roland Leins is a Project Manager at the International Technical Support Organization, San Jose Center. He writes extensively and teaches IBM classes worldwide on all areas of distributed storage management with a focus on Tivoli Storage Manager and Storage Area Networks. Before joining the ITSO in 1999, Roland worked in the EMEA Systems Marketing department as EMEA Marketing Manager for ADSM.

Lloyd Dieter is a Senior Technology Consultant with Synergy, Inc., an IBM Premier Business Partner in Rochester, NY. He has 18 years of experience in the Information Technology field. Before joining Synergy in May of 1998, he worked for Hewlett-Packard Corporation as a Software-Peaked Customer Engineer. He holds a degree in Electrical Engineering Technology, and is also an IBM Certified Advanced Technical Expert (CATE). His areas of expertise include enterprise backup and system recovery on various platforms.

Wim Feyants is an I/T Specialist in Belgium. He has 6 years of experience in the technical support and implementation field, both on Intel and S/390 platforms. His areas of expertise include Tivoli Storage Manager on S/390, Windows NT, and AIX, with 4 years of experience altogether. He holds a degree in Electromechanical Engineering. His previous publications include the IBM Redbooks, *IBM Netfinity and PC Server Technology and Selection Guide,* SG24-4760, and *Netfinity Tape Solutions*, SG24-5218.

Leos Stehlik is an IT Specialist in IBM Czech republic. He has 4 years of experience in the UNIX and Windows NT fields. His areas of expertise include Tivoli Storage Manager on AIX, HP-UX, Windows NT, and OS/390, with 3 years of experience altogether in these fields.

Thanks to the following people for their invaluable contributions to this project:

Michael Hrencecin
IBM Tucson

Karen Orlando
IBM Tucson

Martin Tiernan
IBM Tucson

MIchael Kaczmarski
Tivoli Systems, Tucson

William Scheid
Tivoli Systems, Tucson

Alejandra Vande Zande
Tivoli Systems, Tucson

Avishai Hochberg
Tivoli Systems, San Jose

David Derk
Tivoli Systems, San Jose

Rob Edwards
Tivoli Systems, San Jose

Omar Vargas
Tivoli Systems, San Jose

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Please send us your comments about this or other Redbooks in one of the following ways:

- Fax the evaluation form found in “IBM Redbooks review” on page 299 to the fax number shown on the form.
- Use the online evaluation form found at **ibm.com**/redbooks
- Send your comments in an Internet note to redbook@us.ibm.com

Chapter 1. SAN overview

This chapter provides a brief overview of current SAN technology, addressing the questions:

- What is a SAN?
- Why are SANs needed?
- How are SANs constructed?

1.1 What is a SAN?

A Storage Area Network, or SAN, can be defined as a high-speed network to move data between heterogeneous servers and storage resources.

IBM defines a SAN in the following manner:

A Storage Area Network (SAN) is a dedicated, centrally managed, secure information infrastructure, which enables any-to-any interconnection of servers and storage systems.

A valuable resource for understanding SAN technology is the publication *Fibre Channel: Connection to the Future,* 1-878701-45-0. This publication provides a good non-vendor-specific presentation of the evolving SAN landscape. An excellent overview of IBM's SAN components is provided in the publication *Introduction to Storage Area Network, SAN*, SG24-5470.

1.2 Why are SANs needed?

In the past decade there has been a shift from the older host-centric computing model to the distributed storage model that arises from the client-server computing model. In this client-server model, storage is attached to a particular server, which then may or may not be shared with other servers across a network using communication protocols such as TCP/IP. Storage attached to a single server in this manner is sometimes referred to as server attached storage (SAS).

There are several disadvantages to this model:

- Having storage attached solely to one server tends to lead to “islands” of storage, where not all servers have high-speed access to the storage they may require.

- The communication network, usually a LAN, is typically shared, and is used not only for data movement, but for other tasks, such as user communication. Large data transfers of the type that might occur between servers can lead to significant network traffic problems. While increasing the speed of the network (for example, from 10 Mb/second to 100Mb/second) helps to ease these types of problems temporarily, it is only a stopgap measure.
- Communications network protocols, like IP, are fairly inefficient for moving large blocks of data around.
- While LANs and WANs are very flexible, they are relatively easily disrupted.
- It is relatively easy to monitor or “hijack” traffic traveling on a LAN or WAN. While encryption is in use to some degree, the bulk of traffic in the information technology world is not encrypted.

A Storage Area Network addresses these issues:

- As all servers attached to the SAN can be made to “see” all of the storage available, you no longer have isolated “islands” of storage visible to only one server.
- A SAN is usually dedicated for data traffic. As user communication traffic does not travel on the SAN, network bottleneck problems are alleviated.
- SANs today are usually implemented on Fibre Channel, and typically provide throughput of 100MB/second. This is significantly faster than even Gigabit ethernet.
- SANs today are almost always implemented on Fibre Channel protocols; FCP is much more efficient for moving large quantities of data.
- Fibre Channel is usually implemented on optical fiber, which is stable, flexible, and extremely difficult to “tap” or hijack.

Further information on this topic is available in these publications:

- *Introduction to Storage Area Network, SAN*, SG24-5470
- *Storage Area Networks: Tape Future in Fabrics*, SG24-5474
- *Fibre Channel: Connection to the Future,* 1-878701-45-0

1.3 Fibre Channel topologies

Fibre Channel connections can be grouped into one of three general categories:

- Point-to-point connections
- Loop connections, also known as Fibre Channel Arbitrated Loop (FC-AL)
- Switched fabrics

This section discusses these different Fibre Channel topologies, and how they are used today.

1.3.1 Point-to-point topology

The point-to-point topology is the simplest Fibre Channel configuration to build, and the easiest to administer. Figure 1 shows a simple point-to-point configuration.

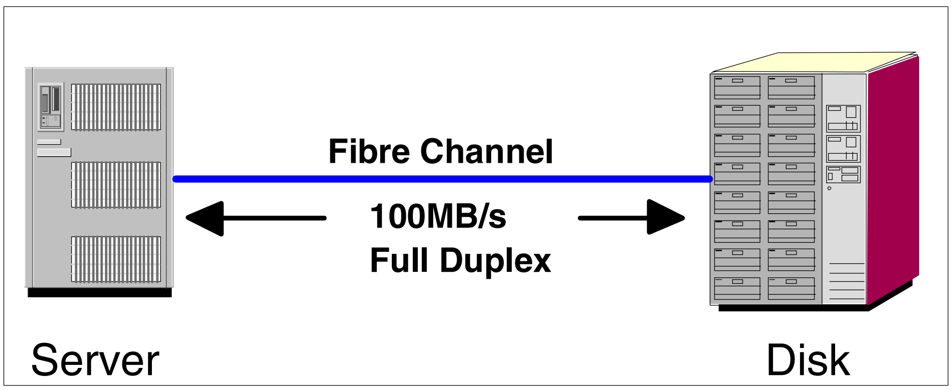

Figure 1. Point-to-point Fibre Channel connection

Point-to-point connections are most frequently used between servers and storage devices, but may also be used for server-to-server communications.

1.3.2 Arbitrated loop topology

Fibre Channel Arbitrated Loop topology, also known as FC-AL, is a fairly low-cost alternative that provides high-speed connectivity for up to 126 devices on a single loop. Figure 2 shows a Fibre Channel loop using a hub.

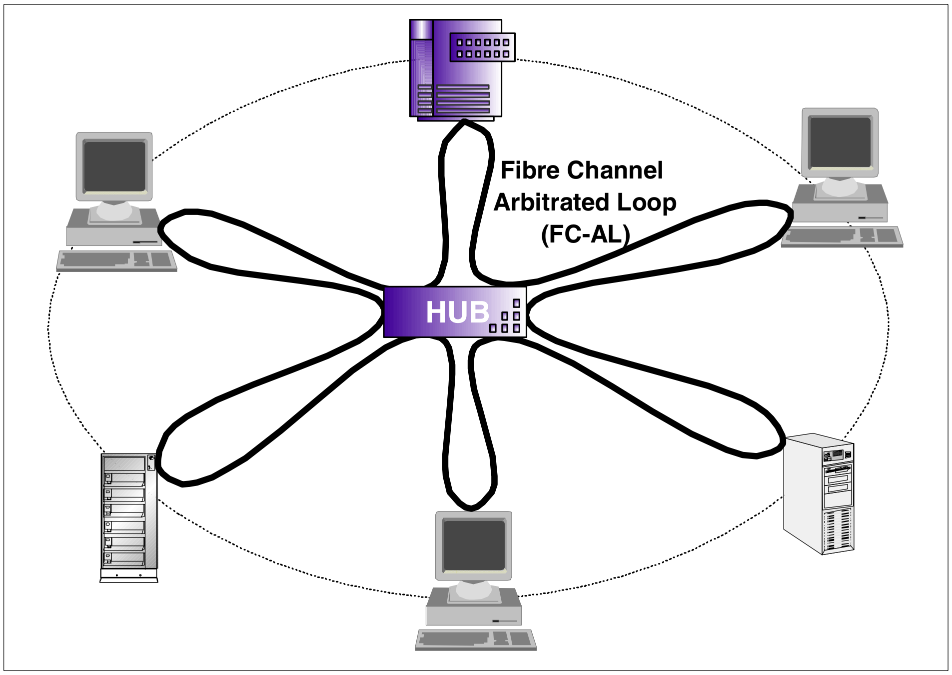

Figure 2. Fibre Channel arbitrated loop (FC-AL)

In an arbitrated loop configuration, all devices attached to the loop must request control of the loop when they wish to send data. This process of requesting control is called *arbitration*. Another characteristic of this type of connection is that all devices share the available bandwidth. Accordingly, as more devices are added to the loop, effective throughput decreases.

In the event that a device loses connectivity to the loop, hubs typically include logic to isolate the node and maintain continuity in the loop.

1.3.3 Switched fabric

Fibre Channel switched fabrics provide the highest speeds and greatest expandability. Reasons for using switches include:

- Increased bandwidth
- Scalable performance
- Increased number of attached devices
- Possibly improving availability through redundancy

A switched fabric may be cascaded or non-cascaded.

Non-cascaded switched fabric

In a non-cascaded fabric, SAN switches are not connected together, but are only connected to servers and storage devices. Figure 3 shows a non-cascaded switched fabric.

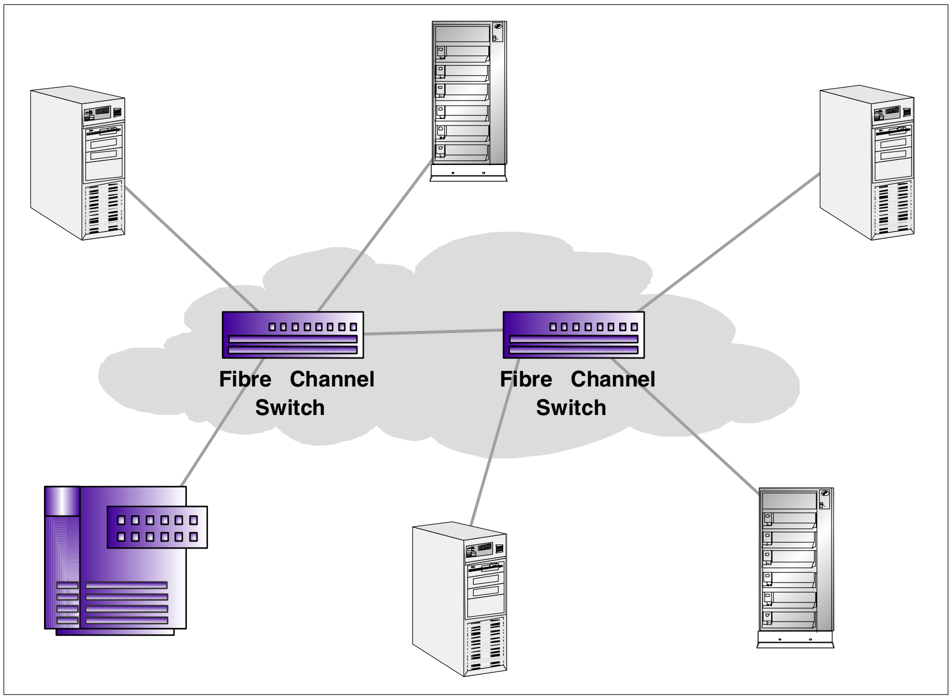

Figure 3. Non-cascaded switched fabric

Note that there are no links between the switches shown in Figure 3. In a non-cascaded fabric, every switch is managed separately.

This type of configuration provides alternate paths in the event that a link or switch fails, and can allow any system to get to any device.

Cascaded switched fabric

Figure 4 shows an example of a cascaded switched fabric.

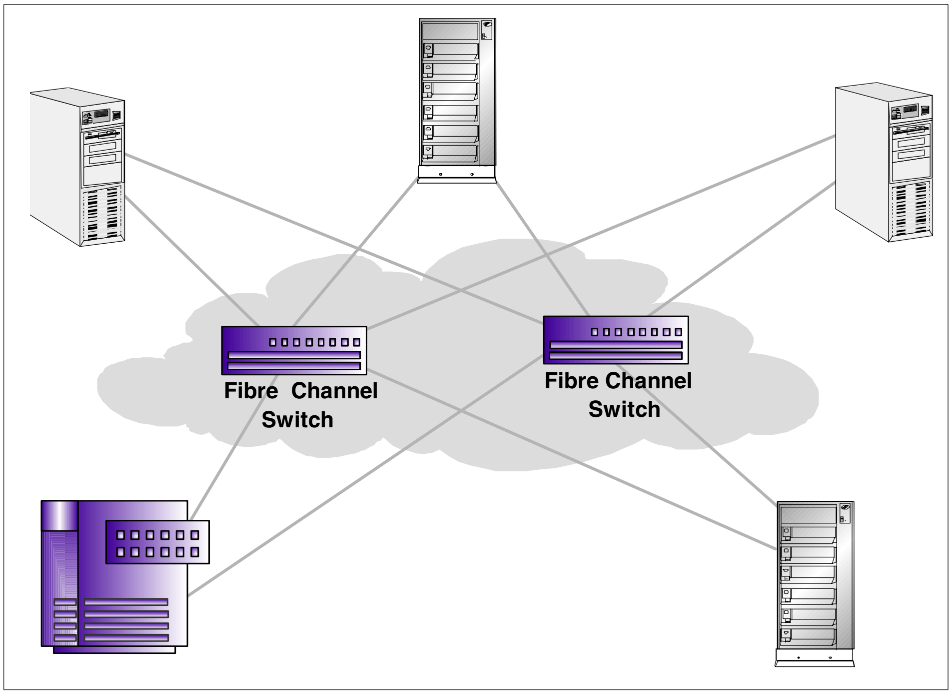

Figure 4. Cascaded switched fabric

In a cascaded configuration, switches are connected together by inter-switch links. Configuration information is shared (cascaded) between switches using these links.

When using a cascaded configuration, multiple switches can be made to function as one large switch. The disadvantage of this type of configuration is that it can be more difficult to configure and manage.

1.4 Overview of SAN components

This section provides an overview of the SAN components which we used in the making of this publication; however, it is not our intention to discuss every available SAN component.

Although a storage area network can be implemented using several different interconnect technologies, including SCSI, almost all SAN implementations today are being built with Fibre Channel. Given this fact, we elected to confine our test environment to one built on Fibre Channel devices.

1.4.1 Cables and connectors

Fibre Channel is supported using both copper cabling and fiber optic cabling. Fiber optic cabling can be one of two types: *multi-mode* or *short-wave* fiber, which is used for shorter distance applications (up to 2 kilometers), and *single-mode* or *long-wave* fiber, which is used for distances up to 10 kilometers. Table 1 outlines IBM's support for the different fiber optic types.

Table 1. IBM's support for the different fiber optic types

Diameter (microns)	Mode	Laser Type	Distance
9	single-mode	long wave	up to 10 km
50	multi-mode	short wave	Up to 500 m
62.5	multi-mode	short wave	up to 175 m

There are also connectors and adapters that allow the interconnection of fiber optic devices with copper-based devices. Refer to *Introduction to Storage Area Network, SAN*, SG24-5470 for further information.

1.4.2 Host Bus Adapters

Host Bus Adapters (HBAs) reside inside the server (host) to be attached to the SAN. They provide the interface from the server to the transmission media (copper or fiber) that the SAN is based on. In the context of this publication, we refer to them as HBAs.

In this publication, we utilized the following Host Bus Adapters:

- Gigabit Fiber Channel Adapter for PCI Bus (Type 4-S) for RS/6000 (feature code 6227)
- QLogic QLA2100F Fibre Channel PCI adapter
- Qlogic QLA2200A Fibre Channel PCI adapter

1.4.3 Gigabit Interface Converters

Gigabit Interface Converter (GBIC) typically reside insides a hub or switch. These can be thought of as an “adapter”, converting the electrical or optical signal from the cable into something the hub or switch can understand. They perform a function for the hub or switch which is similar to what an HBA does for a server.

You must use the correct type of GBIC for the media you are using; in other words, if you are using long-wave single-mode fiber, that is the type of GBIC that you must use also. In our lab setup, we used GBICs for short-wave 50 micron fiber cable.

1.4.4 Hubs

A *SAN hub* performs a very similar function for a SAN that a LAN hub does for a LAN.

A Fibre Channel hub can connect up to 126 devices in a single logical loop. All the devices on this loop share the available bandwidth of the hub. A Fibre Channel hub can be *managed* or *unmanaged*; some of the managed hubs (including IBM's 3534 Managed Hub) have characteristics and capabilities similar to Fibre Channel switches.

For the purposes of this publication, and the scenarios described in our lab, we used the IBM Fibre Channel Hub, model 2103-H07.

1.4.5 Gateways

A *SAN gateway* is a device which is used to interconnect dissimilar networks or devices, and which may or may not perform protocol conversion.

The IBM SAN Data Gateway 2108-G07 which we used in creating this publication is technically a router. It connects between the Fibre Channel connections on one side, and the fast-wide differential connections needed for our 3570 tape library on the other.

1.4.6 Routers

A *SAN router* performs a function very similar to its namesake in the LAN world. The difference is that instead of using a messaging protocol such as TCP/IP, a SAN router uses a storage protocol like FCP.

The IBM SAN Data Gateway Router used in this publication looks and acts very similar to the SAN Data Gateway discussed in the last section; it provides similar, but somewhat more limited capabilities.

At a glance, the differences between the router and the gateway are these:

- The gateway has more Fibre Channel ports: The gateway can accept up to 6 Fibre Channel connections when configured with 3 dual-port PMCs. The router can accept a maximum of one dual-port PMC, providing two short-wave Fibre Channel connections.

- The gateway can accept either short-wave and long-wave connections, depending on which PMC cards are installed.
- The gateway provides more SCSI connections: Four SCSI connections are provided on the gateway, as opposed to two on the router.
- The gateway can provide *virtual private SAN* (VPS) capabilities (optional): With the separately purchased VPS option, you can create virtual private SANs, giving you finer control over your environment than is provided with just zoning.

1.4.7 Switches

Fibre Channel switches are the highest performance devices for connecting large numbers of nodes and storage devices, and providing the highest throughput.

When Fibre Channel switches are connected together in a SAN, the resulting “mesh” of connections is called a fabric. This type of connection method can provide on-demand connection capability from virtually any device to any device. In the event that a single switch does not provide enough connectivity, multiple switches can be connected together; this is referred to as “fanout”. Examples of configurations using switches can be found in Section 1.3.3, “Switched fabric” on page 4.

For the purposes of this document, we used an IBM Fibre Channel Switch, model 2109-S08.

Chapter 2. Introduction to Tivoli Storage Manager

This chapter introduces the Tivoli Storage Management product set, focussing on the Tivoli Storage Manager product. In addition, we will look at specific information on how and why Tivoli Storage Manager can be implemented in a SAN environment.

2.1 Architecture and basic concepts

Tivoli Storage Manager is the core product of the Tivoli Storage Management product set. It provides a solution for distributed data and storage management in an enterprise network environment. Tivoli Storage Manager supports a wide variety of platforms for mobile, small and large systems, and delivers, together with complementary products, many data management functions, such as data protection for file and application data, record retention, space management, and disaster recovery.

This section gives a high-level technical introduction to Tivoli Storage Manager. It positions Tivoli Storage Manager within the Tivoli Storage Management solution, provides an overview of its architecture, the base concepts, the interfaces, and supported environments, and shows Tivoli Storage Manager's interaction with other products of the Tivoli Storage Management product set.

2.1.1 The Tivoli Storage Management solution

In today's connected world, data has become the key asset of companies and one of its most important competitive differentiating factors. Temporary inaccessibility or the complete loss of data has a huge financial impact, and can drive companies out of business. The inability to manage data can have a negative impact on a company's profitability and limit their ability to grow. Storing, protecting, and managing data growth has become one of the major challenges of today's businesses.

Based on this requirement, Tivoli defined its *Information Integrity Initiative:* "The Tivoli Storage Management Initiative provides an end-to-end software management solution with proven methodologies to help customers link storage management policies with key business practices, to enable them to use information to DRIVE business, rather than simply support it."

Figure 5 shows the structure of the Tivoli Storage Management Solution and how it fits into the Tivoli Enterprise.

Tivoli Storage Management consists of four major solution components:

- Enterprise protection
- Application protection
- SAN management
- Workgroup protection

Figure 5. Tivoli Storage Management and Tivoli Enterprise

Enterprise protection implements an enterprise-wide solution for data protection, disaster recovery, space management and record retention. It covers all types of heterogeneous system platforms, starting from mobile systems up to large scale enterprise servers, and supports all types of storage resources, locally attached as well as network or SAN attached storage. Flexible storage management policies and powerful automation features support business needs by eliminating labor-intensive and cost-intensive manual storage management tasks.

Strategic business applications are typically complex collections of interdependent components from both commercial and proprietary software, and span desktop, distributed, and mainframe computing environments. Application protection is concerned with data availability, performance, and the ability to recover and integrate application data management into enterprise data protection.

Storage Area Network (SAN) management is concerned with the efficient management of the Fibre Channel based SAN environment. Physical connectivity mapping, switch zoning, performance monitoring, error monitoring, and predictive capacity planning are among the most important features.

Workgroup protection provides a reliable, easy to use, backup, recovery, and disaster recovery solution for stand-alone mobile, desktop and small server systems. It is targeted to small and medium businesses (under 800 clients) as well as any enterprise with large numbers of remote, stand-alone servers.

Combined with Tivoli Enterprise, Tivoli Storage Management becomes an integrated management suite that transforms Information Technology into a strategic business resource.

2.1.2 Tivoli Storage Manager

Tivoli Storage Manager is the core product of the Tivoli Storage Management product set. It provides a solution for distributed data and storage management in an enterprise network environment. It is the next generation of the product that was originally released by IBM as ADSTAR Distributed Storage Manager (ADSM). Tivoli Storage Manager protects and manages data on more than 30 operating platforms, covering mobile, desktop and server systems over the entire distributed world. It is integrated with hundreds of storage devices as well as LAN, WAN, and emerging SAN infrastructures.

The base function provided by Tivoli Storage Manager and its complementary products can be divided into two main categories: data protection and storage management.

- The **Data Protection** category includes:
  - **Operational backup and restore of data:** The *backup process* creates a copy of the data to protect against the operational loss or destruction of file or application data. The customer defines how often to back up (frequency) and how many numbers of copies (versions) to hold. The *restore process* places the backup copy of the data back into a customer-designated system or workstation.

- **Disaster recovery:** All activities to organize, manage, and automate the recovery process from a major loss of IT infrastructure and data across the enterprise. This includes processes to move data offsite into a secure vault location, to rebuild IT infrastructure, and to reload data successfully in an acceptable time frame.

• **The Storage Resource Management** category includes:

  - **Vital record retention, archive and retrieval:** The *archive process* creates a copy of a file or a set of files representing an end point of a process for long term storage. Files can remain on the local storage media or can be deleted. The customer controls how long (retention period) an archive copy is to be retained. The *retrieval process* locates the copies within the archival storage and places them back into a customer-designated system or workstation.
  - **Hierarchical space management (HSM):** This process provides the automatic and transparent movement of operational data from the user system disk space to a central storage repository. If the user accesses this data, it is dynamically and transparently restored to the client storage.

The solution is network based, which means that these functions are available to the whole network environment. All the functions can be automated to run in a 24x7 lights-out environment. Administration costs are minimized by centralization of all of the management of Tivoli Storage Manager components.

2.1.2.1 Tivoli Storage Manager architecture

Tivoli Storage Manager is implemented as a client server software application, consisting of a Tivoli Storage Manager server software component, Tivoli Storage Manager Backup/Archive client, the Tivoli Storage Manager storage agent, and other complementary Tivoli and vendor software products. Figure 6 shows the main components of Tivoli Storage Manager.

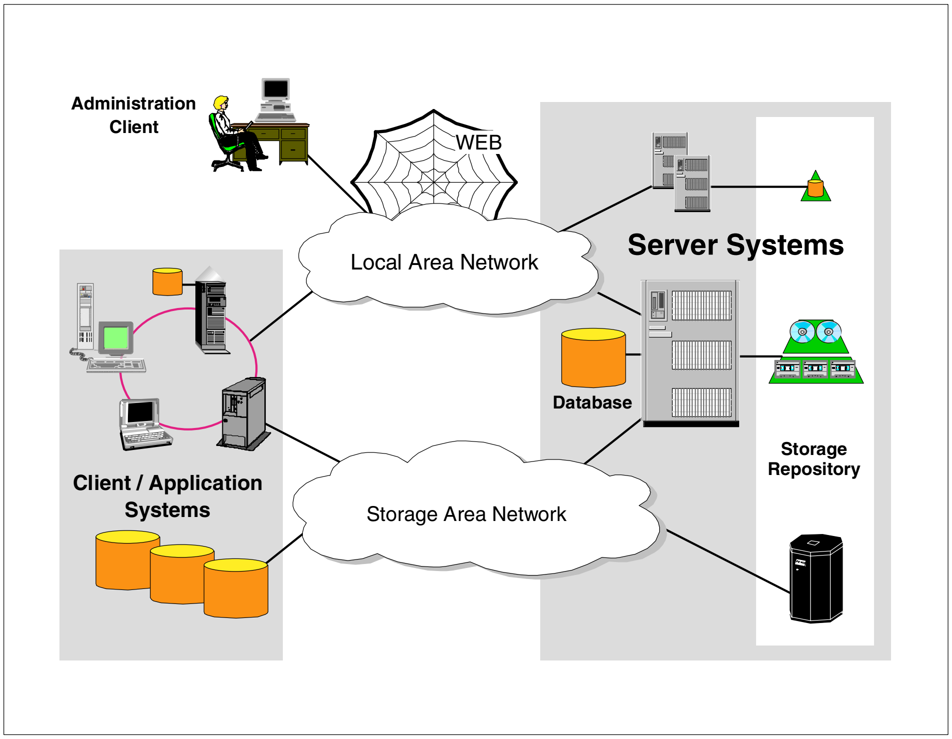

Figure 6. Tivoli Storage Manager architecture

The *Tivoli Storage Manager server software* builds the data management backbone by managing the storage hardware, providing a secure environment, providing automation, reporting and monitoring functions, and implementing the storage management policies and by storing all object inventory information in the *Tivoli Storage Manager database*. The *Tivoli Storage Manager client software, Tivoli Storage Manager storage agent software,* and *complementary products* implement data management functions like data backup and recovery, archival, hierarchical space management or disaster recovery.

The client software can run on different systems, including laptop computers, PCs, workstations, or server systems. The client and server software can also be installed on the same system for a local backup solution. The storage agent software in conjunction with the server software enables the implementation of LAN-free backup solutions exploiting SAN infrastructure. It is also possible to define server hierarchies or multiple peer-to-peer servers in order to provide a multi-layer, centrally managed storage management solution or an electronic vaulting solution.

Management of servers and client node definitions is done using the *Tivoli Storage Manager Administrative* client. This client connects to one or more servers through a Web interface, or using a command line interface.

Tivoli Storage Manager server

One of the principal architectural components of the Tivoli Storage Manager server software is its in-built relational database. The storage manager database was especially designed for the task of managing data, and it implements zero-touch administration. All policy information, logging, authentication and security, media management, and object inventory is managed through this database. Most of the fields are externalized through Tivoli Storage Manager high level administration commands, SQL SELECT statements, or for reporting purposes, by using an ODBC driver.

For storing the managed data, the Tivoli Storage Manager server uses the *storage repository.* The storage repository is designed from any combination of disk, optical, tape, or robotic storage devices, which are locally connected to the server system or which are accessible through a SAN. Exploiting the SAN technology, the server software has features implemented to dynamically share SAN connected automated tape library systems among multiple Tivoli Storage Manager server systems.

The server software provides built-in drivers for more than 300 different device types from every major manufacturer. It is also able to utilize operating device drivers and external library manager software such as the Windows 2000 Removable Storage Manager (RSM).

Within the storage repository the devices can operate stand-alone or can be linked together to form one or more storage hierarchies. The *storage hierarchy* is not limited in the number of levels and can also span over multiple servers using so-called *virtual volumes*. See the Section “Storage and device concepts” on page 27 for storage management function defined on the storage repository.

Tivoli Storage Manager storage agent

To support LAN-free backup solutions exploiting SAN infrastructures, starting with Tivoli Storage Manager Version 4.1 a special *Tivoli Storage Manager storage agent* software component has been introduced on the Tivoli Storage Manager client system. The storage agent software dynamically shares SAN connected and server controlled tape library resources with the Tivoli Storage Manager server software, and it has the ability to write directly to server-owned tape storage media in a format which is consistent with that used by the server.

The storage agent software receives data objects from the Tivoli Storage Manager data management application programming interface (as discussed in the Section “Tivoli Storage Manager externalized interfaces” on page 20) and communicates with the Tivoli Storage Manager server software over the LAN using TCP/IP to exchange control information and metadata about the objects being backed up, but the data movement utilizes the LAN-free path over the SAN to write directly to the tape storage media. So the data movement is off-loaded from the LAN and from the Tivoli Storage Manager server processor for greater scalability.

Tivoli Storage Manager backup/archive client

Data management functions are implemented using Tivoli Storage Manager client software and complementary Tivoli and non-Tivoli products, which work together with the Tivoli Storage Manager server backbone product.

The *Tivoli Storage Manager backup/archive client*, included with the server program provides the operational backup and archival function. The software implements the patented *progressive backup methodology, adaptive sub-file backup technology* and unique record retention methods as described in the Section “Backup and archival concepts” on page 25.

All version 3.7 and above backup/archive clients are implemented as multi-session clients, which means that they are able to exploit the multi-threading capabilities of modern operating systems. This enables the running of backup and archive operations in parallel to maximize the throughput to the server system. Full exploitation of multi-threading on the client also requires a 3.7 or higher Tivoli Storage Manager server.

Depending on the client platform, the backup/archive client may provide a graphical, command line, or Web user interface (see Figure 7). Many platforms provide all three interfaces. The command line interface is helpful for experienced users and allows generation of backup or restore scripts for scheduled execution. The graphical interface is designed for ease of use for the end user for ad hoc backups and restores. The Web client is especially useful for those clients, such as NetWare, where no native GUI is available, or for performing remote backup/restore operations, for example, in a helpdesk environment.

Some UNIX based Version 3.7 clients use a new plug-in architecture to implement an image backup feature for raw device backup. This allows you to backup and recover data that is not stored in file systems or supported database applications. It also provides an additional method to make point-in-time backups of entire file systems as single objects and recover them in conjunction with data backed up by using the progressive backup methodology.

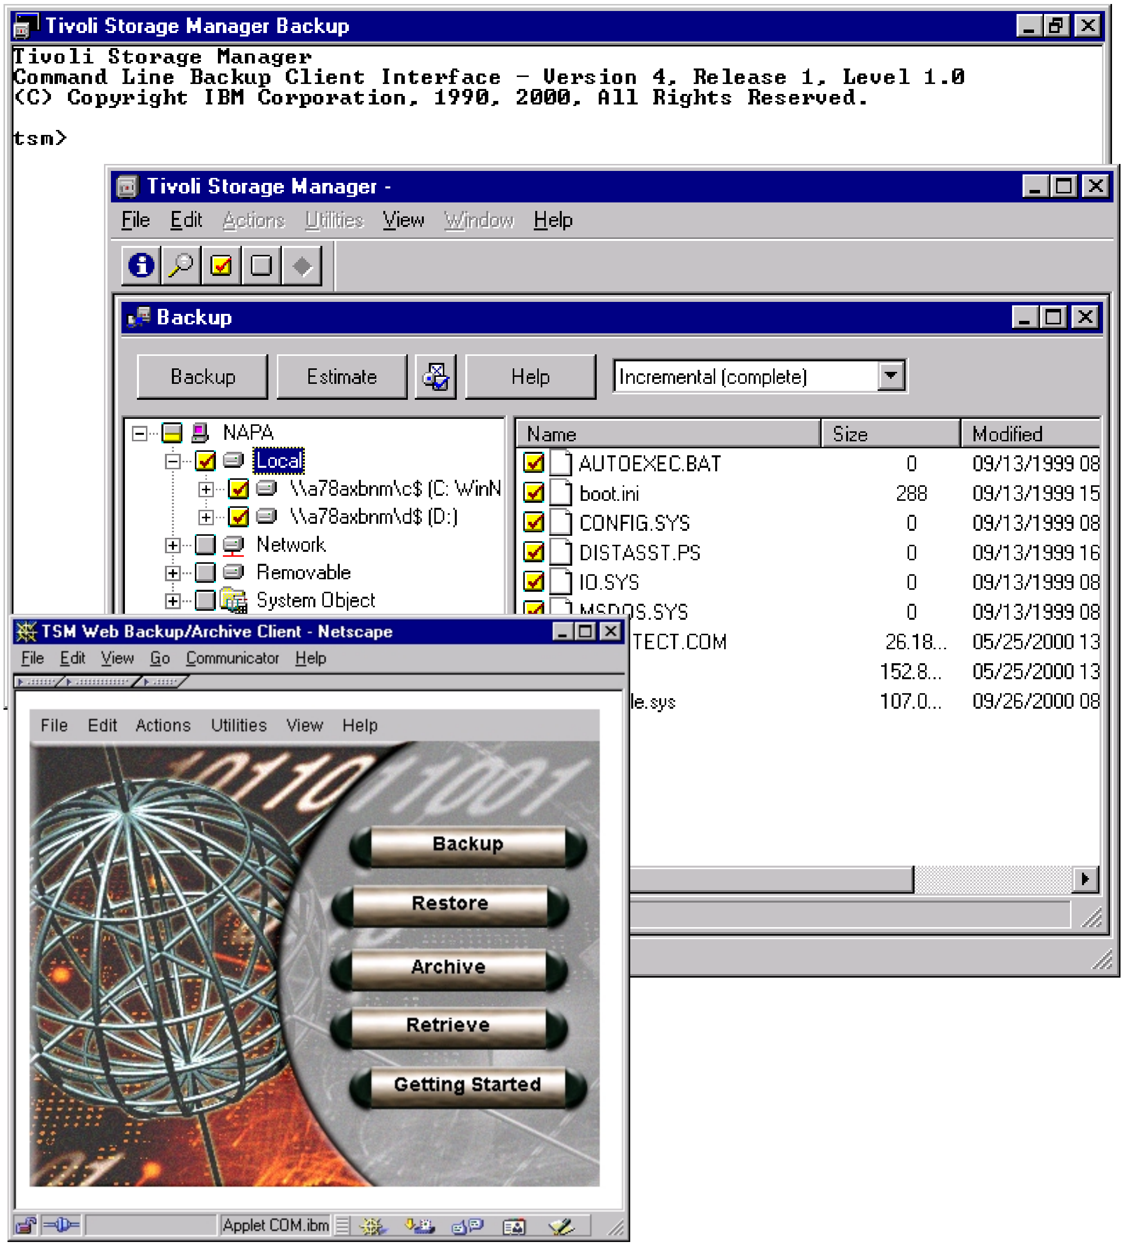

Figure 7. Backup Archive client user interfaces

Tivoli Storage Manager administration
For the central administration of one or more Tivoli Storage Manager instances, and the whole data management environment, Tivoli Storage Manager provides command line or Java-based administration interfaces (see Figure 8), also called *administration clients*.

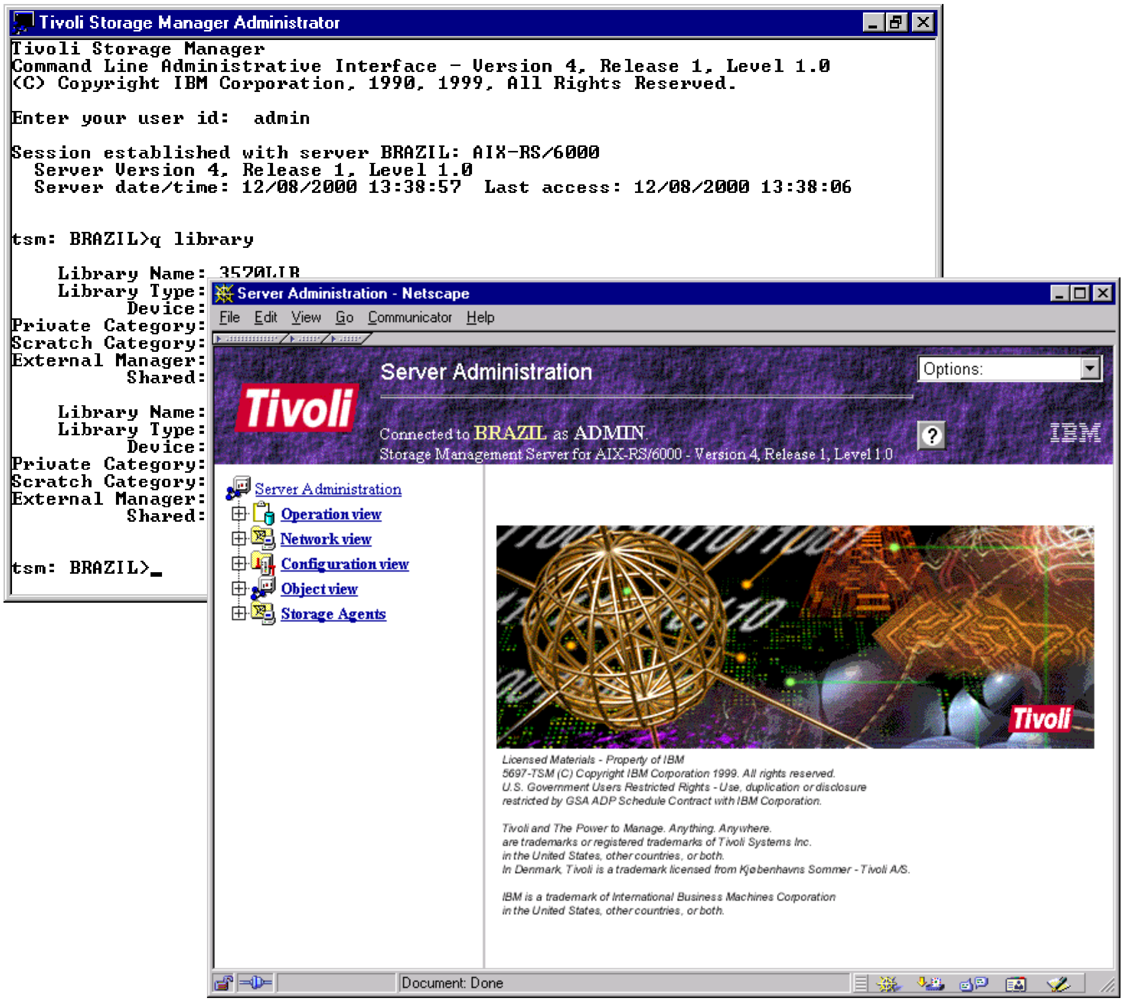

Figure 8. Tivoli Storage Manager administration interfaces

Using the unique *enterprise administration* feature it is possible to configure, monitor and manage all server and client instances from one administrative interface, known as the *enterprise console*. Enterprise administration includes:

- Enterprise configuration
- Administrative command routing
- Central event logging functions

The enterprise configuration allows certain server configurations to be defined centrally by an administrator and then propagated to other servers on a scheduled base. This simplifies the configuration and management of multiple Tivoli Storage Manager servers in an enterprise significantly.

Administrative command routing allows administrators to issue commands from one server and route them to other target servers or groups of servers. The commands are executed on the target servers, and the command output is returned and formatted on the server where the command was issued.

In an enterprise environment with multiple Tivoli Storage Manager servers, client and server events can be logged to a central management server through server-to-server communications, thereby enabling centralized event management and automation.

Tivoli Storage Manager externalized interfaces

Tivoli Storage Manager provides a data management *application programming interface* (API), which can be used to implement *application clients* to integrate popular business applications, such as databases or groupware applications, into the Tivoli Storage Management solution. The API is also published to allow customers or vendors to implement specialist clients for special data management needs or non-standard computing environments. In general we distinguish between *Tivoli Data Protection for application* software products and the API exploitation through vendor applications.

Tivoli Data Protection for Applications are separate program products delivered by Tivoli to connect business applications to the Tivoli Storage Manager data management API. Such applications, for example, Oracle, Lotus Notes and Domino, Microsoft Exchange, and Microsoft SQL server, have their own storage management interfaces which are used to interface. Another area for Tivoli Data Protection program products is to exploit the special features of *intelligent disk subsystem*, for example, IBM's Enterprise Storage Server or EMC's Symmetrix within a Tivoli Storage Manager based storage management solution. For more information, see Section 2.1.4, "Tivoli Data Protection for applications and hardware integration" on page 35.

On the other hand, many vendor applications also exploit the Tivoli Storage Manager data management API by integrating it directly into their software product to implement new data management functions, or to provide backup and archival functionality on additional system platforms. Some examples are IBM's CommonStore for R/3 data archival, IBM's BRMS/400 to provide an AS/400 backup solution, and SSSI's Archive Backup Client for OpenVMS data backup and recovery.

In addition to the externalized interfaces to the server database as described in the Section "Tivoli Storage Manager server" on page 16, Tivoli Storage Manager offers multiple interfaces for event logging, reporting and monitoring the data management environment. In general, activities of the Tivoli Storage Manager server and client are logged in the server database and can be sent for reporting and monitoring purposes to external *event receivers* using *event filter* mechanism. Some potential event receivers are the Tivoli Enterprise Console (TE/C), SNMP based systems management software packages, the Windows NT event log, and user written applications.

To integrate the Tivoli Storage Manager storage management capability with external library management applications, Tivoli Storage Manager offers an *external library manager interface*. Using this interface it is possible to integrate the Tivoli Storage Manager server into third-party storage management environments.

Tivoli Storage Manager supported environments

Tivoli Storage Manager server and client software is available on many different operating system platforms and can exploit different communication protocols. Figure 9 shows an overview of the supported environment.

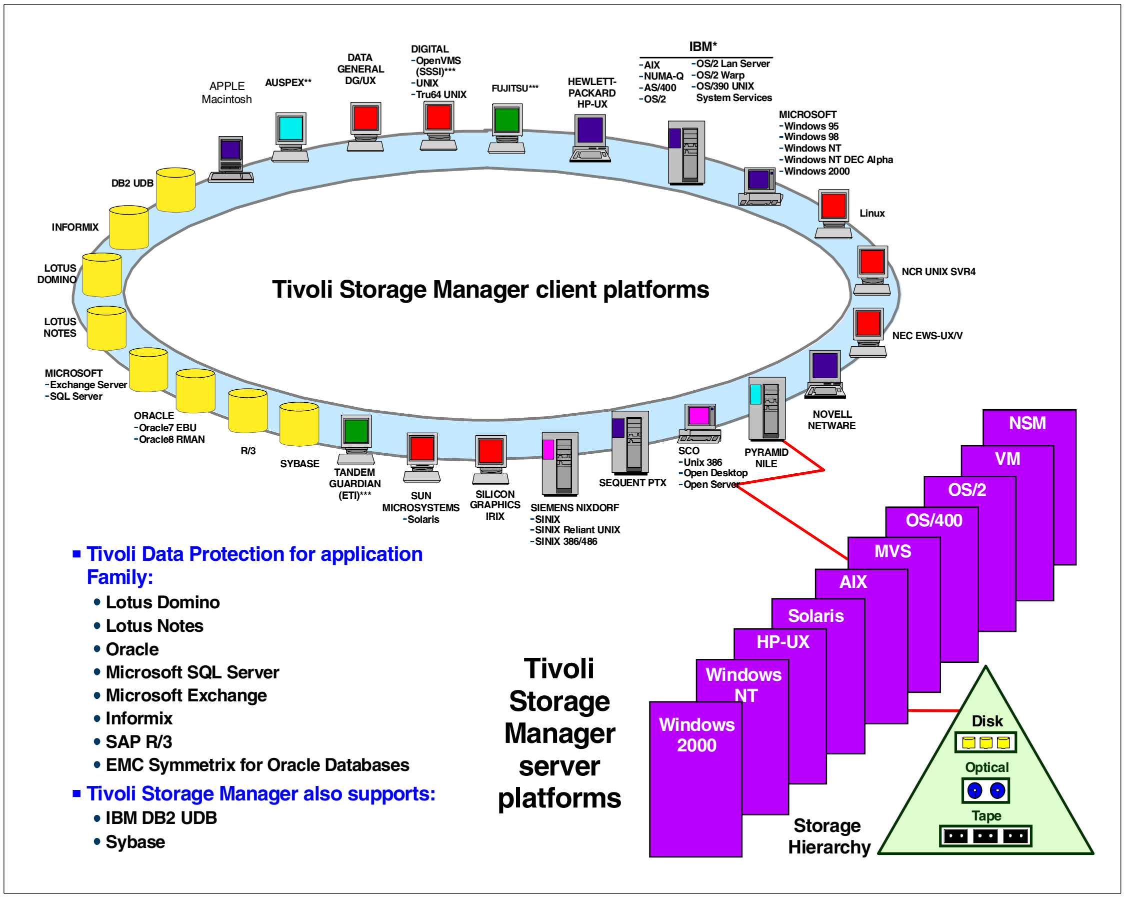

Figure 9. Tivoli Storage Manager supported environments

The Tivoli Storage Manager server software runs on nine different operating platforms plus the IBM Network Storage Manager (NSM). At this time, the software is available in the most recent Version 4.1 on the platforms, as shown in Table 2, illustrating server platforms, operating system, and currently available Tivoli Storage Manager versions.

Table 2. Tivoli Storage Manager most recent server platforms

Server platforms	Operating system level	Server version
AIX	4.3.2., 4.3.3 (4.3.3 required for SAN implementations)	4.1
HP-UX	11.0	4.1
MVS/ESA OS/390	Version 5 Release 2.2 or higher Version 2 Release 1 or higher	4.1
Sun Solaris	2.6, 7	4.1
Windows NT	Workstation 4.0, Server 4.0 (SP 3/4/5/6/6a) (SP5 or higher required for SAN implementations)	4.1
Windows 2000		4.1
OS/400	V4R3 or higher	3.1.2
VM/ESA	2.2,2.3,2.4	3.1
OS/2	WARP 4.0 or higher WARP Server 4.0	2.1

Table 3 and Table 4 provide an overview of all available Version 3.7 and 4.1 clients at the time of publishing this paper. Other operating system clients (available at pre-3.7 levels) may be updated in the future. Check the product information on the Tivoli Storage Manager home page for the latest complete client availability information:

`http://www.tivoli.com/products/index/storage_mgr/`

There are several variations of UNIX clients. Table 3 details the UNIX clients and the operating system levels that are supported.

Table 3. Version 3.7 and 4.1 UNIX clients

Client Platforms	Version	**Operating System**
AIX	4.1	4.3.2, 4.3.3
HP-UX	4.1	11.0
Sun Solaris	4.1	2.6, 7
Linux	4.1	SUSE 6.3, RedHat 6.1, Turbo Linux 6.0, Calder Linux 2.3
OS/390 UNIX System Services	4.1	Version 2 Release 8 or Version 2 Release 9
SCO UnixWare	4.1	7.01
Tru64 UNIX	4.1	4.0d or higher
SGI IRIX	4.1	6.5.x
IBM NUMA-Q (formerly Sequent)	4.1	DYNIX/ptx 4.5.1 or higher

There are also different PC clients available. Table 4 details the PC systems and the operating systems that are supported as clients.

Table 4. Version 3.7 and 4.1 PC clients

PC Clients Platforms	Version	**Operating Systems**
Novell NetWare	4.1	4.11, 4.20, 5.0, 5.1
Microsoft Windows (Intel)	4.1	NT 4.0, Win 95, Win 98, Win 2000
Microsoft Windows (Alpha)	3.7	NT 4.0
Apple Macintosh for PowerPC	4.1	MacOS 8.x, MacOS 9.x
IBM OS/2	3.7	Warp 4.0 or later

2.1.2.2 Base concepts

This section gives a high level introduction to the base data and storage management paradigms used by Tivoli Storage Manager to implement its functionality. We will cover data protection or backup, record retention or archival, storage management, policy, and security.

Backup and archival concepts

Backup, in Tivoli Storage Manager terms, means creating an additional copy of a data object to be used for recovery. A data object can be a file, a part of a file, a directory or a user defined data object like a database or database table. The backup version of this data object is stored separately in the Tivoli Storage Manager server storage repository. Potentially, you can make several backup versions of the data, each version at a different point-in-time. These versions are closely tied together and related to the original object as a group of backups.

If the original data object is corrupted or lost on the client system, r*estore* is the process of sending a backup version of the data from the server back to the client. The most current version of the data is normally restored, but you can choose to restore from any of the existing backup versions. The number of backup versions is controlled by server definitions. Old backup versions may be automatically deleted as new versions are created. You may also delete them after a certain period of time.

For file level based backup the main difference from many other backup applications is that Tivoli Storage Manager uses the *progressive backup methodology*. As shown in Figure 10, after the first necessarily full backup, Tivoli Storage Manager then operates with incremental backups only. In consequence, only those files that have changed since the last backup will be backed up.

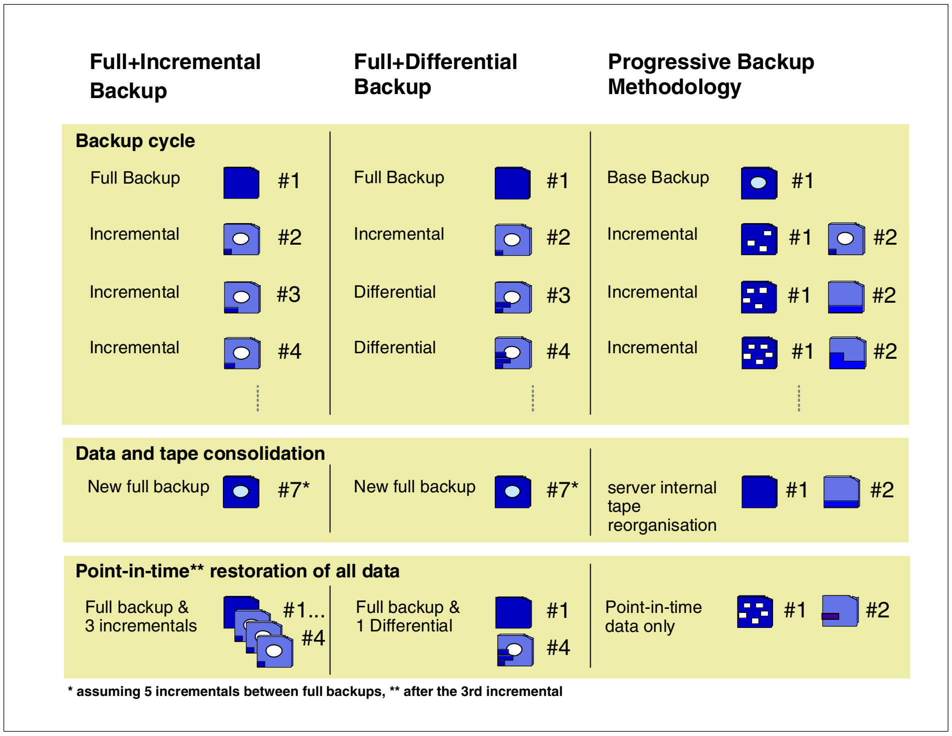

Figure 10. Progressive backup methodology versus other backup schemes

The reorganization of the physical storage media to store each client's data physically together on a small number of media—in order to provide faster access in the case of a complete system recovery—is done transparently to the client, and is completely automated on the server using metadata information stored in the server database.

Tivoli Storage Manager's file level progressive backup methodology, in comparison with other methods like Full+Incremental or Full+Differential backup schemes, prevents unnecessary backups of unchanged data to reduce and consolidate the recovery tape-set. It offers also a more efficient use of storage resources by not storing redundant data and a faster recovery by not restoring multiple versions of the same file.

Tivoli Storage Manager's *adaptive sub-file backup technology* implements another powerful method to furthermore reduce the amount of data transferred from the client to the server system. The features enables the backup-archive client (Web client, command line and GUI) to back up only the changed portion of a file, either on byte or on block level, instead of transferring the whole file to the server every time. Especially for mobile or remote client systems, this feature helps significantly to overcome bandwidth limitations of the network link.

At any point in time Tivoli Storage Manager allows the creation of a complete set of client files (*backup set*) on the server system using the most recent backup versions stored in the server storage repository. These backup sets can be used to retain a snapshot of all client files for a longer period of time (*Instant Archive*) or for LAN-free recovery of a client system by copying this backup set onto portable media and restoring them locally (*Rapid Recovery*).

File *Archive* with Tivoli Storage Manager means creating a copy of a file as a separate object in the storage repository to be retained for a specific period of time. Typically you would use this function to create an additional copy of data to be saved for historical purposes. Vital records (data that must be kept for legal or other business reasons) are likely candidates for the archive process. You can specify to delete the original copy of the data on the source system once the archive copy is created on the server. Therefore, you can use *archive* to make additional space available on the Tivoli Storage Manager client system. However, *archive* should not be thought of as a complete space management function, because transparent automatic recall is not available.

You can access archived data by using *retrieve* to return it to the Tivoli Storage Manager client, if the data is needed at some future time. To locate the archived data within the storage repository, Tivoli Storage Manager allows you to add a description to the data and to form *archive packages* of related files. You can then use this description to search the server database for matching packages, to determine which data to retrieve.

Therefore, the difference between *backup* and *archive* is that backup creates and controls multiple backup versions that are directly attached to the original file; whereas archive creates an additional file that is normally kept for a specific period of time, as in the case of vital records.

Storage and device concepts

All client data that is managed by Tivoli Storage Manager is stored in the Tivoli Storage Manager storage repository. This repository can consist of different storage devices, such as disk, tape, or optical devices, and controlled by the Tivoli Storage Manager server. To do so, Tivoli Storage Manager uses its own model of storage to view, classify, and control these storage devices, and to implement its storage management functionality (see Figure 11).

The main difference between the storage management approach of Tivoli Storage Manager and other commonly used systems is that the Tivoli Storage Manager concentrates on managing *data objects* instead of managing and controlling backup tapes. Data objects can be sub-file components, files, directories or raw logical volumes that are backed up from the client systems; they can be objects like tables or records from database applications, or simply a block of data that a client system wants to store on the server storage.

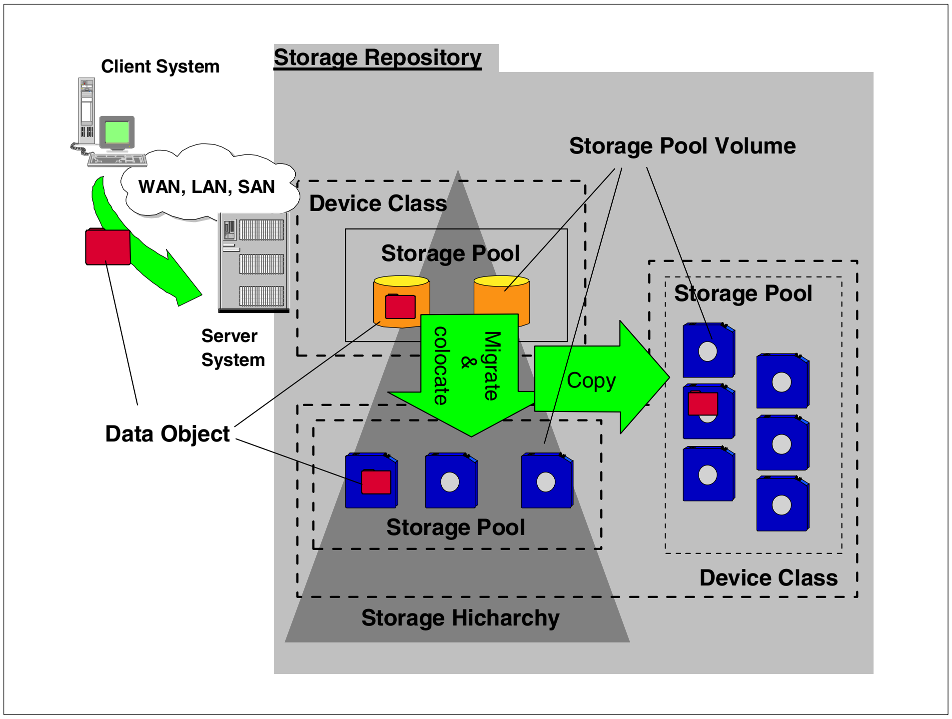

Figure 11. Tivoli Storage Manager storage management concept

To store these data objects on storage devices and to implement storage management functions, Tivoli Storage Manager has defined some logical entities to classify the available storage resources. Most important is the logical entity called a *storage pool.* A storage pool describes a storage resource for one single type of media, such as a disk partition or a set of tape cartridges. Storage pools are the place where data objects are stored.

A storage pool is built up from one or more *storage pool volumes*. For example, in the case of a tape storage pool, this would be a single physical tape cartridge. To describe how Tivoli Storage Manager can access those physical volumes to place the data objects on them, Tivoli Storage Manager has another logical entity called a *device class.* A device class is connected to a storage pool and specifies how volumes of this storage pool can be accessed.

Tivoli Storage Manager organizes storage pools in one or more hierarchical structures. This *storage hierarchy* can span multiple server instances, and is used to implement management functions to *migrate* data objects automatically — completely transparent to the client — from one storage hierarchy level to another; or in other words, from one storage device to another. This function may be used, for example, to cache backup data (for performance reasons) onto a Tivoli Storage Manager server disk space before moving the data to tape cartridges. The actual location of all data objects is automatically tracked within the server database.

Tivoli Storage Manager has implemented additional storage management functions for moving data objects from one storage volume to another. As discussed in the previous section, Tivoli Storage Manager uses the progressive backup methodology to backup files to the Tivoli Storage Manager storage repository. The reorganization of the data and storage media for fast recovery happens completely within the server. For this purpose, Tivoli Storage Manager has implemented functions to *relocate* data objects from one volume to another and to *collocate* data objects that belong together, either at the client system level or at the data group level.

Another important storage management function implemented within the Tivoli Storage Manager server is the ability to *copy* data objects asynchronously and to store them in different storage pools or on different storage devices, either locally at the same server system or remotely on another server system. It is especially important for disaster recovery reasons to have — in the event of losing any storage media or the whole storage repository — a second copy of data available somewhere in a secure place. This function is fully transparent to the client, and can be performed automatically within the Tivoli Storage Manager server.

Policy concepts

A data storage management environment consists of three basic types of resources: client systems, rules, and data. The client systems contain the data to be managed, and the rules specify how the management must occur; for example, in the case of backup, how many versions should be kept, where they should be stored, and so on.

Tivoli Storage Manager policies define the relationships between these three resources. Figure 12 illustrates this policy relationship. Depending on your actual needs for managing your enterprise data, these policies can be very simple or very complex.

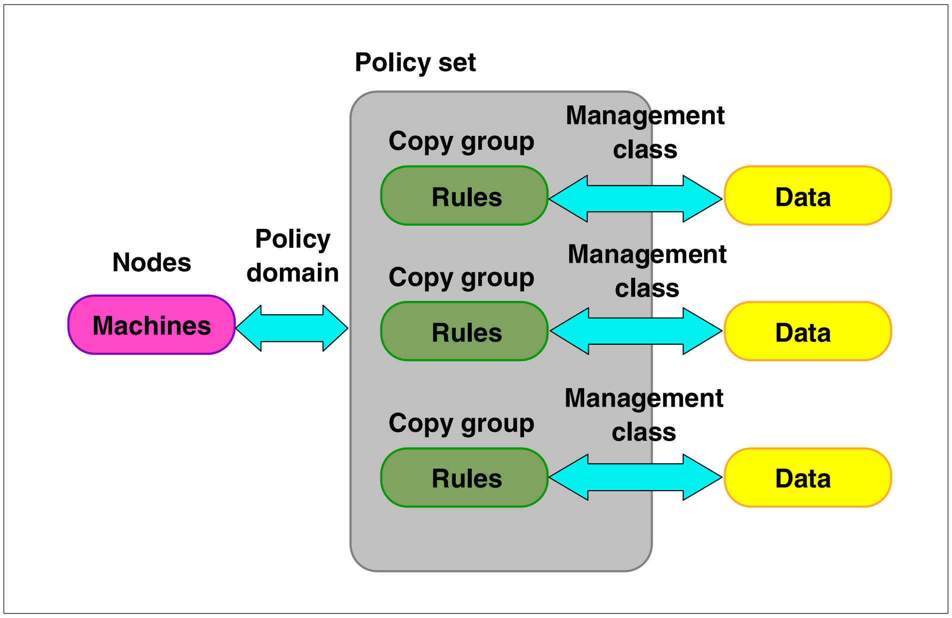

Figure 12. Policy relations and resources

Tivoli Storage Manager has certain logical entities that group and organize the storage resources and define relationships between them. Client systems, or nodes in Tivoli Storage Manager terminology, are grouped together with other nodes with common storage management requirements, into a *policy domain*.

The policy domain links the nodes to a *policy set*, a collection of storage management rules for different storage management activities. A policy set consists of one or more *management classes*. A management class contains the rule descriptions called *copy groups*, and links these to the data objects to be managed. A copy group is the place where all the storage management parameters, such as number of stored copies, retention period, storage media, and so on, are defined. When the data is linked to particular rules, it is said to be bound to the management class that contains those rules.

Another way to look at the components that make up a policy is to consider them in the hierarchical fashion in which they are defined. That is to say, consider the policy domain containing the policy set, the policy set containing the management classes, and the management classes containing the copy groups and the storage management parameters.

Security concepts

Because the storage repository of Tivoli Storage Manager is the place where all the data of an enterprise are stored and managed, security is a very vital aspect for Tivoli Storage Manager. To ensure that data can only be accessed from the owning client or an authorized party, Tivoli Storage Manager implements, for authentication purposes, a *mutual suspicion algorithm*, which is similar to the methods used by Kerberos authentication.

Whenever a client (backup/archive or administrative) wants to communicate with the server, an authentication has to take place. This authentication contains *both-sides* verification, which means that the client has to authenticate itself to the server, and the server has to authenticate itself to the client.

To do this, all clients have a password, which is stored at the server side as well as at the client side. In the authentication dialog these passwords are used to encrypt the communication. The passwords are not sent over the network, to prevent hackers from intercepting them. A communication session will be established only if both sides are able to decrypt the dialog. If the communication has ended, or if a time-out period without activity is passed, the session will be automatically terminated and a new authentication will be necessary.

In mobile computing environments, files are often sent to the Tivoli Storage Manager server system using a modem connection, and so they are exposed to the security hazards of public telephone lines. The Tivoli Storage Manager backup/archive client implements (in addition to the end-point security concept outlined above) a *data encryption* function, which allows for encrypting data before it is sent to the Tivoli Storage Manager server, and which protects the data while it is being transferred to the server and also while it resides in the storage repository.

2.1.3 Tivoli Storage Manager complementary products

Tivoli Storage Manager complementary products use the Tivoli Storage Manager server software package as a backbone product to implement additional data and storage management functions. The following section introduces Tivoli Space Manager for hierarchical space management, Tivoli

Disaster Recovery Manager as an enterprise-wide solution for disaster recovery and Tivoli Decision Support for Storage Management Analysis for a comprehensive reporting and monitoring solution to plan the growth and collect vital management information for an efficient enterprise data management deployment.

2.1.3.1 Tivoli Space Manager

Tivoli Space Manager uses the framework services of Tivoli Storage Manager in combination with the industry standard *Data Management Application Programming Interface (DMAPI)* to deliver a fully integrated solution for open systems *Hierarchical Space Management (HSM)*. Tivoli Space Manager provides an HSM client, which interfaces with DMAPI and implements the functionality outlined in Figure 13.

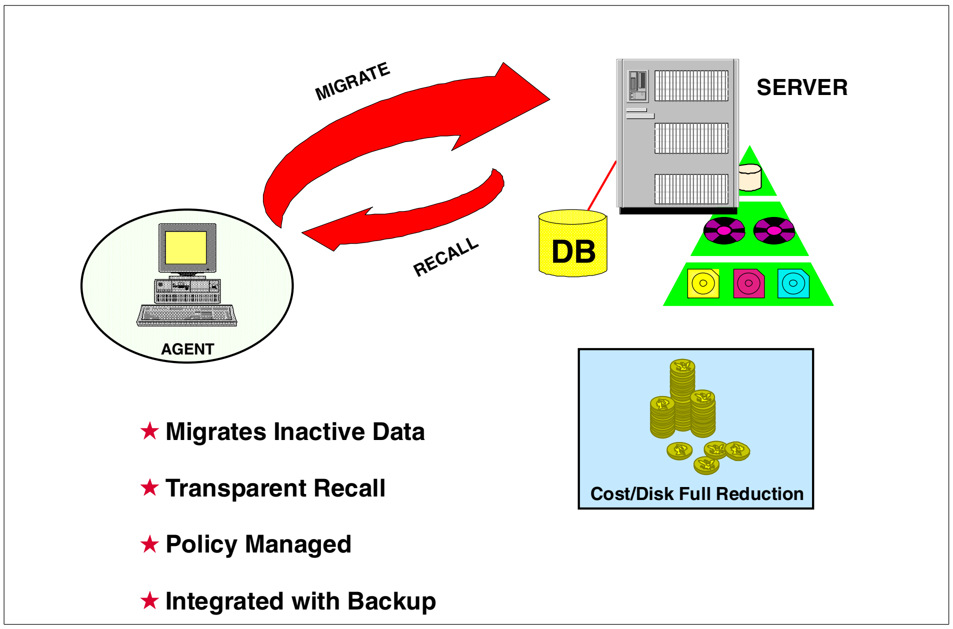

Figure 13. Hierarchical storage management

Tivoli Space Manager maximizes usage of existing storage resources by transparently migrating data from client hard drives to the Tivoli Storage Manager storage repository based on size and age criteria. When the migrated data is accessed, Tivoli Space Manager transparently recalls it back onto the local disk.

The migration of files and the management of migrated files is controlled by policies. However, user controlled migration and recall is also possible.

Tivoli Space Manager's HSM function is fully integrated with Tivoli Storage Manager operational backup. It is possible to specify not to migrate a file until it has a backup version in the server storage repository. If a file is migrated and then a backup is done the next day, Tivoli Storage Manager copies the file from the Space Migration pool to the backup pool, instead of requiring a recall to the client system to back it up again, which would cause multiple transfers of the same file across the network.

2.1.3.2 Tivoli Disaster Recovery Manager

Tivoli Disaster Recovery Manager coordinates and automates the process of recovering from a disaster. It integrates with Tivoli Storage Manager and the rest of the Tivoli data management portfolio to provide for offsite media management, automated restore and managed recovery. It complements the already implemented robust protection features of Tivoli Storage Manager and automates many already facilitated protection functions.

Tivoli Disaster Recovery Manager automatically captures information required to recover the Tivoli Storage Manager server after a disaster. It assists in preparing a plan that allows recovery in the most expedient manner. This *disaster recovery plan* contains information, scripts, and procedures needed to automate server restoration, and helps ensure quick recovery of your data after a disaster.

Tivoli Disaster Recovery Manager also manages and tracks the movement of off-site media to reduce the time required to recover in the event of a disaster. It is able to track media that are stored on-site, in-transit, or off-site in a vault, no matter whether it is a manual or electronic vault, so your data can be easily located if disaster strikes.

Client recovery information can also be captured by Tivoli Disaster Recovery Manager to assist with identifying what clients need to be recovered, in what order, and what is required to recover them.

2.1.3.3 Tivoli Decision Support for Storage Management Analysis

Tivoli Decision Support for Storage Management Analysis uses the framework services of Tivoli Decision Support to deliver important decision-making information about your enterprise data management deployment.

Tivoli Decision Support is a stand-alone product that provides a ready-to-use view into the wealth of data gathered by Tivoli enterprise products. The product consolidates this data from Tivoli products and transforms it into accessible IT business-relevant information. This information, presented in a variety of graphical formats can be viewed interactively (slice, dice, drill down, drill through) and posted on a URL. Tivoli Decision Support provides insight and the ability to better answer IT business-relevant questions. Tivoli Decision Support is available on Windows NT.

The architecture and the information flow is shown in Figure 14.

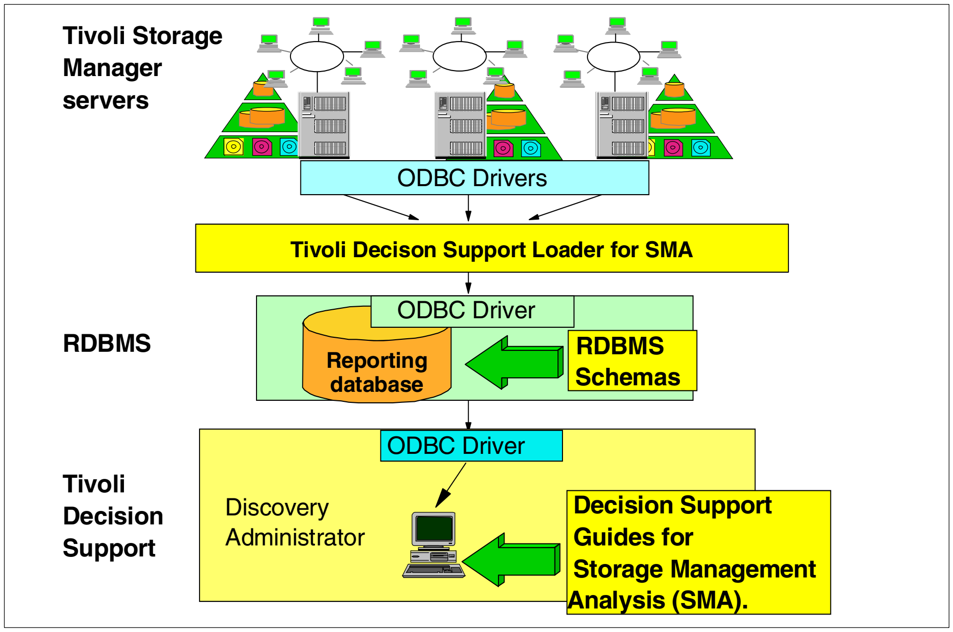

Figure 14. Tivoli Decision Support for Storage Management Analysis

The Tivoli Decision Support Discovery Guides are a set of best-practices guides provided for various applications. To use these guides Tivoli Decision Support has to be installed and available. The Tivoli Decision Support for Storage Management Analysis is the guide used to produce the following analyses:

- Storage Event Analysis
- Storage Performance Analysis

The information used by the guide is obtained directly from the Tivoli Storage Manager server with the use of the ODBC interface. The information is then transferred to a relational database, another requirement for Tivoli Decision Support for Storage Management Analysis. The databases supported for the information feed are DB2, MS SQL, and Oracle. The database can reside on the same system as Tivoli Storage Manager or Tivoli Decision Support or on a separate system. The database is used to for queries to generate the Tivoli Decision Support reports.

2.1.4 Tivoli Data Protection for applications and hardware integration

Tivoli Data Protection for applications is a group of solutions integrated to Tivoli Storage Manager, which protect data used by business applications. They are interface programs between a storage management API provided by the vendor application, and the Tivoli Storage Manager data management API. Typical applications providing such interfaces are databases and groupware applications, such as Lotus Notes or Microsoft Exchange.

Figure 15 shows Tivoli Data Protection for Lotus Domino, a typical example of architecture and data flow of a Tivoli Data Protection for Application solution.

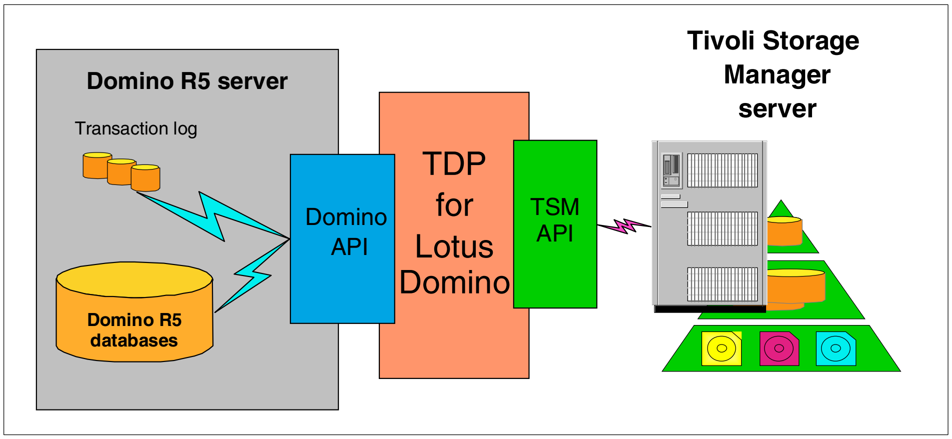

Figure 15. Tivoli Data Protection for Lotus Domino

The function of the Tivoli Data Protection for applications solutions is to receive application backup and restore requests and to translate them into Tivoli Storage Manager backups and restores. The activity is always initiated from the application. This means that backups or restores can also be done while the application is on line. The TDP application client therefore implements the intelligence and function of the backup interface provided by the application itself.

However, in 7x24 production environments, even the ability to do an online backup of that application data is not an option. For example, intelligent disk subsystems (such as IBM's Enterprise Storage Server or EMC's Symmetrix) implement features which allow for creating a copy of all the application data on different physical volumes that are accessible by an alternate (backup) system, with nearly no interruptions. Tivoli Data Protection for Hardware Integration program products interface with these features and integrate them into a complete Tivoli Storage Manager data management solution.

Table 5 shows the available Tivoli Data Protection for Applications and hardware integration solutions and the platforms, operating system level and application level they cover at the time of publishing this publication.

Table 5. Tivoli Data Protection for application and hardware integration solution

TDP solution	**Application level**	**Operating system**	**Operating system level**
TDP for Lotus Notes V2.1.9	4.5.x, 4.6.0, 4.6.1, 4.6.3 - 4.6.6	AIX NT	4.2.x, 4.3.x 4.0 SP4
TDP for Lotus Domino V1.1.1	5.0.1	AIX NT Windows2000 Solaris OS/390 USS	4.3.1, 4.3.2 4.0 SP4 2.6 or 7 2.6, 2.7
TDP for MS Exchange V1.1.1	5.0, 5.5	NT Windows2000	4.0 SP4
TDP for MS SQL Server V1.1.2	6.5, 7.1	NT Windows2000	4.0 SP4
TDP for Informix V4.1.0	IDS 7, UDO 9	AIX Sun Solaris	4.3.2, 4.3.3 2.6, 7
TDP for Oracle V2.1.10	7.3.4, 8.0.X, 8.1.5, 8.1.6	AIX NT Sun Solaris HP-UX Windows2000	4.3.1, 4.3.2, 4.3.3 4.0 2.6, 7 (64bit) 11.0
TDP for R/3 V3.1	3.0C to F, 3.1 up to Z, 4.0 up to 4.6C	AIX HP-UX Sun Solaris Tru64 UNIX NT (Intel) Windows 2000 LINUX	4.3 11.0 7 4.0 4.0 RedHat 6.1 EE
TDP for EMC Symmetrix for Oracle V1.1.0	8.1.5	Sun Solaris	7

For the latest information on Tivoli Data Protection solutions, check the following Tivoli Storage Management Web page:

`http://www.tivoli.com/products/solutions/storage/`

2.1.4.1 Tivoli Data Protection for Workgroups

Tivoli Data Protection for Workgroups provides a complete disaster and bare metal recovery solution for Windows NT servers.

It uses Stac's unique *Object Replication Technology* to ensure backup of all disk contents, such as partition information, boot records, FAT and NTFS partitions, while allowing full read and write access to the system during backup. It can backup entire Windows NT servers or volumes at near maximum device speed to local tape or disk drives.

Object Replication provides an option to the two traditional backup methods. Instead of working at the sector level like image backup or at the file level like file-based backup, this new technology works at the logical "object" level. Objects are defined as each area of the server, including disk partition tables, boot volume, security information, system volume and user data volumes. All data on servers is treated as a single logical unit rather than as separate files, and all data can be copied at once. Object Replication creates an exact copy of the entire server by saving it logically, block by block.

Tivoli Data Protection for Workgroups also provides the function to create a set of potable diskettes that may be used in the event of hard drive failure or other disaster.

The restore operation can be performed using the Windows File Explorer for single file restores by simply dragging and dropping files from the mounted backup medium to their original location; or for more complex restores, by utilizing a Tivoli Data Protection for Workgroups administrative tool. In the event of a complete disk failure, the set of bootable diskettes can be used to perform a bare-metal restore of the entire system.

The product can be either used stand-alone or integrated in a Tivoli Storage Manager environment for the following functions:

- Message logging to a Tivoli Storage Manager server, using the Tivoli Storage Manager API client.
- Automate operations with the scheduler component provided by the Tivoli Storage Manager backup-archive client.

2.2 How to use a SAN environment for backups

This topic discusses advantages a Tivoli Storage Manager implementation can get from SAN attached tape libraries, and eventually from disk storage. These advantages can be grouped into three main areas:

- Availability
- Performance
- Efficiency

2.2.1 Availability

When discussing the availability of a Tivoli Storage Manager, it is important that all separate components are covered. The server can be split into the following three areas:

- The Tivoli Storage Manager server itself, with its database and storage pool volumes.
- The tape storage and SAN fabric.
- System availability, which embodies all activities related to protecting the server operating system and application software.

Each of the solutions discussed in the following sections will refer to one of these areas.

2.2.1.1 Remote tape library connection

When implementing a Tivoli Storage Manager solution, one of the most important issues is server protection. The server database and its storage pool volumes need to be protected against corruption or loss. This is generally done by performing periodic database and storage pool backups to tape volumes, which are then moved to an offsite location. This removal of volumes, and the sometimes required recall can lead to increased complexity and cost of the Tivoli Storage Manager implementation.

Using the capacity of single-mode or long wave fiber connections, a SAN attached tape library can be placed off-site, but remain connected to the Tivoli Storage Manager server. This means that the actual physical checkout of volumes is no longer required, thus reducing cost and complexity. Another advantage is that, in case of primary volume unavailability, the copy pool volumes are instantly available.

Figure 16 shows a possible setup for this solution.

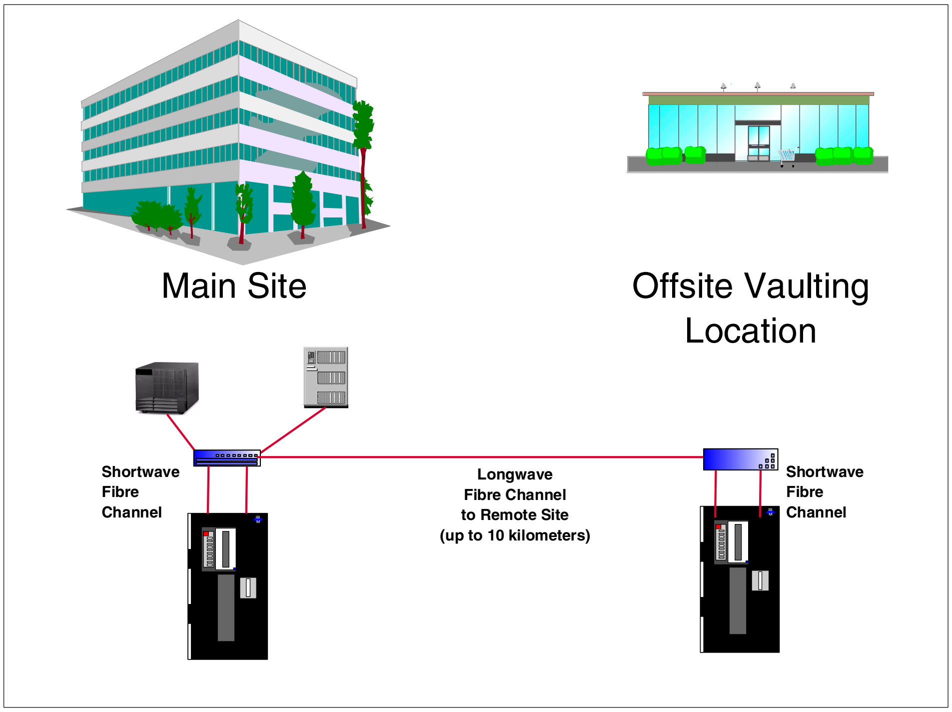

Figure 16. Remotely connected tape library

Locally attached storage, both tape and disk, serves as primary storage space for the server. Database and storage pool backups are redirected to a SAN attached library. Making use of a switch and hub, both equipped with a single-mode fiber GBIC, the distance that can be obtained equals 10 kilometers. This means that the SAN attached library can be located in another building or secure location. In the remote location, the SCSI library can be connected to the SAN using a gateway.

Please refer to Section 8.1, “Remote tape implementations” on page 235 for more information.

2.2.1.2 Fabric fail-over

SAN fabrics are often described as high availability solutions. This statement derives from the fact that when creating a switched fabric, using multiple redundant paths, a switch failure can be handled by the other switches.

Figure 17 shows such a meshed topology. Looking at this configuration, we can see that when a link or a switch fails, other links exists, and other switches can take over the work of the failing switch. Furthermore, this recovery will be automatic.

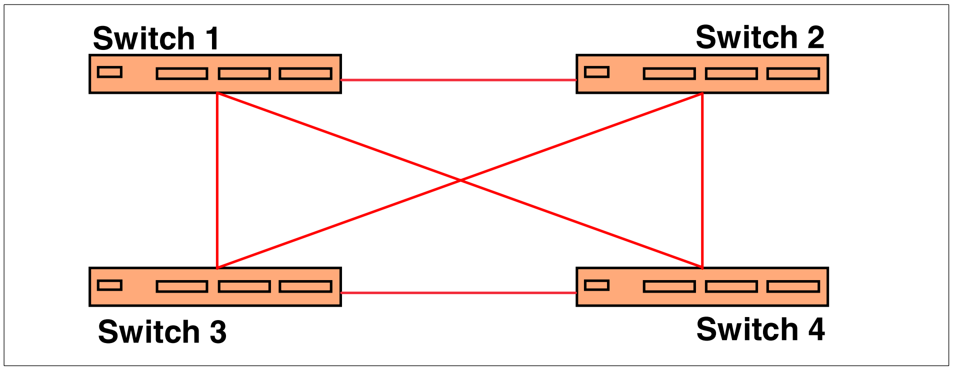

Figure 17. Meshed switch topology

However, caution should be taken when porting this high-availability concept to users of the fabric, and other fabric components. In the case of Tivoli Storage Manager for example, recovery is not automatic. The reason for this is that Tivoli Storage Manager uses the SAN in combination with the SAN Data Gateway, and sees the tape devices as real SCSI devices, rather than native Fibre Channel devices. As we will explain in Chapter 8, “Availability considerations in a SAN” on page 235, recovery of the Tivoli Storage Manager servers after a SAN failure is rather complex, and can be time consuming.

2.2.2 Performance

One of the major advantages and goals of a SAN implementation is the performance gain obtained compared to the classic network based solutions. The reason for these gains are the higher throughput of the SAN fabric and the exclusive character of the fabric (less users than a the normal network). The concept is that data is moved exclusively over the SAN, while the network only servers to transfer metadata or control data between the client and the server. Currently, two implementations of this concept are available using the Tivoli Storage Management product suite. These are the *LAN-free* and the *serverless* or *server-free* solutions.

2.2.2.1 LAN-free implementations

In the Tivoli Storage Manager implementation, the LAN-free backup solution (as shown in Figure 18) has, as its main purpose, the movement data from locally attached client storage to SAN attached Tivoli Storage Manager server tape storage pools, without moving the data over the LAN network. The actual movement and placement of data is under control of the Tivoli Storage Manager server, which communicates with the *Storage Agent* component on the client. This communication flows over the classic LAN network.

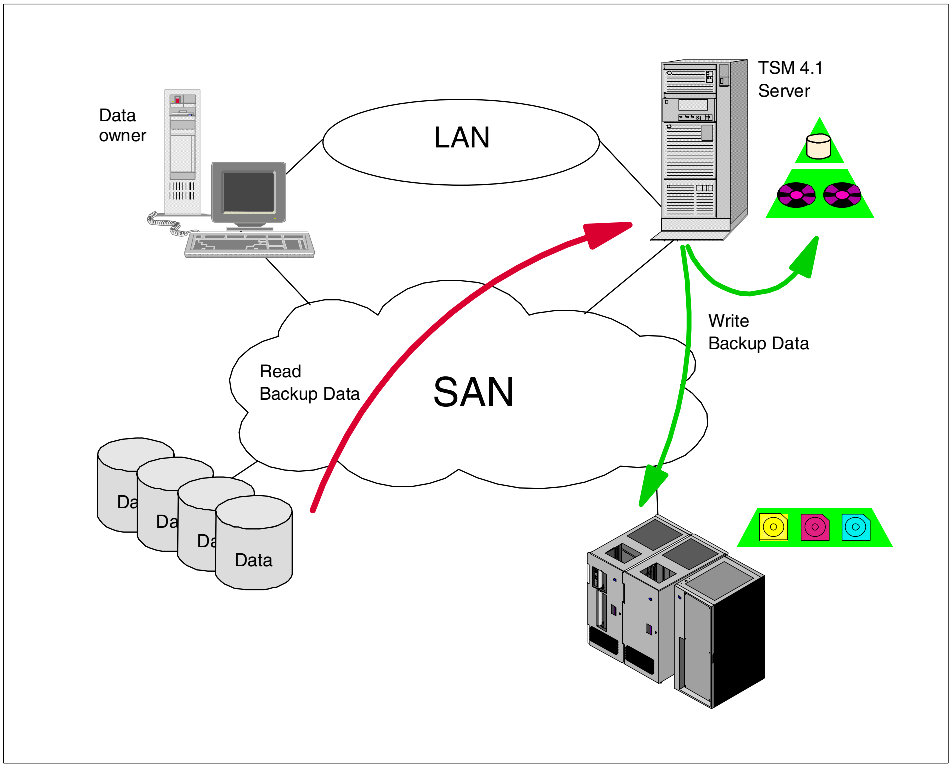

Figure 18. LAN-free backup

A commonly used alternative for this solution is to have a Tivoli Storage Manager server locally installed on the client, with a locally attached SCSI tape library. The disadvantage of such a implementation is that the management of the Tivoli Storage Manager servers becomes more complicated. This is even more important when more than one client are implemented in that way.

Also, the cost of the different tape libraries is important. Another major problem is that the backed up data resides on the same machine as the original data, and the Tivoli Storage Manager server's database. This means that server protection and recovery becomes very important. Also, the processing power, memory, and disk capacity requirements increase.

Currently, the following restrictions apply for LAN-free backup implementations:

- Tivoli Data Protection for MS Exchange and Tivoli Data Protection for SAP R/3 are the only supported clients.
- The Tivoli Storage Manager server and client must be at version 4.1.
- The backup destination can only be a tape storage pool. No disk devices are supported.

For more information on setting up and using a LAN-free environment, please refer to Chapter 5, “LAN-free client data transfer” on page 125.

2.2.2.2 Server-free implementations

Server-free or serverless implementations serve the same main purpose as LAN-free backups, namely to move data over the SAN rather than the LAN. Two major differences exist, however. First, the data resides on SAN attached disk volumes. Second, the owner of the data, which will probably be a file or application server, does not participate in the backup and restore process, hence the name *server-free*. The entire backup process is done by a Tivoli Storage Manager client residing on the same system as the Tivoli Storage Manager server.

Figure 19 shows a possible setup of such a configuration. The data moves to the Tivoli Storage Manager server storage pools using a local Tivoli Storage Manager client (local to the Tivoli Storage Manager server). The storage pool can consist of disk volumes and tape volumes. These tape volumes can belong to a SAN attached tape device. However, this is not a requirement.

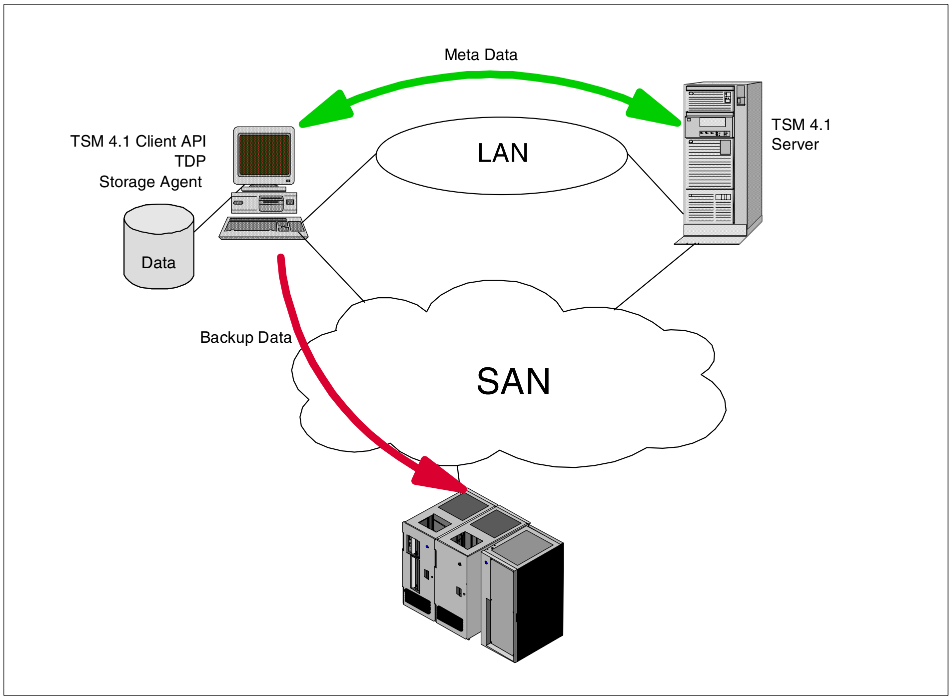

Figure 19. Server free backup

Currently, this implementation can be created using the Tivoli SANergy product. For practical implementation scenarios, please refer to Chapter 7, "Serverless backup using TSM and SANergy" on page 221.

2.2.3 Efficiency

By sharing libraries among several Tivoli Storage Manager servers, it is clear that we will obtain higher tape drive usage, than in cases where every server has its own library. This can also add a cost reduction factor. The reason for this is that you might reduce the total overall number of required tape drives. Before looking at how we can improve the drive usage, we must decide which type of library sharing to use. There are two basic types of library sharing:

- **Split library:** The library is attached to two or more servers using different SCSI connections. Logically, the library must be split into *partitions*, each with their own drives and library slot ranges. This means that each server will only use a small portion of the library, and is unable to access other drives than the ones that where defined to him. When this setup is done using two host machines, it is often referred to as a *twin-tail* configuration. In this case, you are not sharing the drives, but only the library robotics.

- **Library sharing:** As defined in the Tivoli Storage Manager *SCSI Library Sharing* implementation, library sharing means that all hosts attached to the library over a SAN fabric can share the entire library. This means that they have access to all library slots, and all library drives. Tape volumes remain private, meaning that each server only sees the volumes that have been assigned to him. However, if the server needs new tape volumes, he can obtain them from a common scratch tape pool.

2.2.3.1 How can library drive sharing improve the efficiency?

Tivoli Storage Manager servers using non-shared tape libraries typically have a bad utilization rate for the tape devices. The reason for this is that the required number of tape devices is usually given by peak requirements, rather than the required average need. An example is a situation where the requirement is that a restore must be possible, without disrupting server maintenance processes. This would automatically require three tape devices (two for server processes, for example, storage pool backups, and one for the restore session). If multiple concurrent restore sessions are required, this number increases even more. Another situation could be that multiple concurrent backup session require direct tape drive connection, meaning that you would need at least one drive for each of these sessions.

When sharing a library among different servers, the total number of tape devices can be reduced by using this shared capability. The reason for this is that usage becomes more distributed over time, and that peak requirements approach the average usage. For example, if you have two servers, each with the requirement that server maintenance operations and restore session can run at the same time. In a non-shared environment, this would mean that you would need at least three drives per library. In a shared environment however, a first reduction of this number can be obtained by reserving only one drive for restore. This would mean that you could work with only 5 drives.

A further reduction might be possible by phasing the server maintenance processes, in such a way that they do not run at the same moment. As a result, the two drives required for these operations could be shared among the servers. The resulting number of tape devices for the library would then be 3 drives, instead of the original 6. In most cases however, the above calculation is not that simple. In the following section, we will look into library sizing in more detail.

2.2.3.2 Shared library sizing

When implementing a shared tape library, it is important to size the library correctly. By sizing, we mean the storage capacity and the number of tape drives. For the storage capacity, the sizing does not differ from stand-alone implementations. The only difference is that after determining the storage requirements for each server, you will need to add these figures to arrive to the total storage requirement. The number of tape drives that you need becomes more difficult however. We will look at two methods to determine the required number of devices.

Determining the number of drives: method 1

The first and easiest approach to this calculation is to take two tape devices per server that share the library, and a pool of tape devices for backup and restore requirements. The two drives per server are intended to be reserved for server maintenance processes, such as backups of the server database, storage pool backup, migration and reclamation. The required number of tape devices in the backup and restore pool, can be calculated based on the following information:

1. How many *concurrent* backup sessions require tape mounts? This will give you the base number of devices required for the backup sessions. It is always a good practice to minimize this number, using large enough disk storage pools.
2. How many *concurrent* restore sessions will you serve? The answer to this question provides you with the number of tape devices required for restore operations.
3. Do restore and backup sessions occur at the same moment? If yes, add the numbers obtained in steps 1 and 2. If no, the maximum of both numbers should be taken.

The number obtained in the above steps indicates the number of tape drives that are required for the backup and restore pool. Added to the two drives per server required for maintenance, this gives you an initial estimation of the number of tape drives needed.

If the purpose is to minimize the number of drives, the number obtained above can be refined by checking that the drives required for backup and restore sessions cannot be obtained from the server maintenance device pool. The following points can help in determining this:

4. Do backup processes run at the same moment that server maintenance processes? If yes, no sharing of the server maintenance tape drives and shared pool devices is possible. If no, the number of server maintenance devices can be deducted from the number of tape devices required for backup processes (see step 1).
5. Can you interrupt server maintenance processes for restore sessions? Care must be taken when answering this question: in non-shared library environments, the chance that a restore process actually concurs with a server process is small. And if it does, there is a reasonable possibility to delay the server process until the restore operation has finished. In a shared environment however, the chance of restore sessions concurring with server maintenance increases (due to fact that multiple servers share the same library device). If you can interrupt server maintenance processes, deduct the number of server maintenance tape devices from the number obtained in step 2.
6. Do restore and backup sessions occur at the same moment? If yes, add the numbers obtained in steps 4 and 6. If no, the maximum of both numbers should be taken.

Determining the number of drives: method 2

A second method to calculate the required number of tape drives for the shared tape library is to analyze usage of multiple libraries in a current environment. This method is suited for situations where you are migrating from an existing multiple library environment to a single shared one.

Start by checking when and how many tape devices are used for each server that will share a library. Next, these figures should be combined into one time scale. This will show tape requirements in function of the moment in time they occur. Looking at this, there will probably be peak usage moments. The next step would be to see if some processes on some servers can be shifted in time, spreading the usage over a larger period of time, thus reducing the peak requirements.

The following example explains the above procedure in a bit more detail. Figure 20 shows the initial number of tape mounts in three individual libraries. The mounts are classified by the time they are occurring, and the process that is using the tape volumes when they are mounted. Volume mounts for restore operations are not shown. The reason for this is that they are unpredictable, where all other mounts should be rather consistent over a certain time period. This means that after determining the number of drives required, the number needed for restore operations should be added.

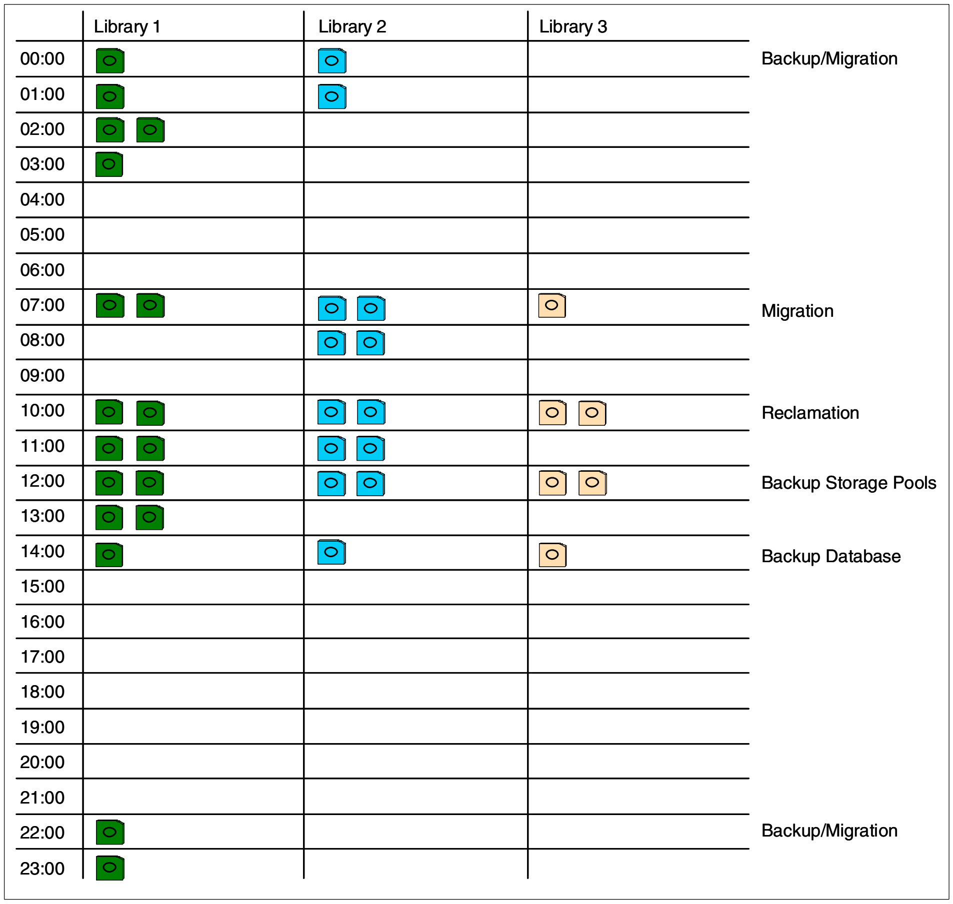

Figure 20. Tape mounts in individual libraries

Next, consolidate these figures, so that they represent the number of mounts required in one shared library. Figure 21 shows the result of this operation. If we keep the processes as shown, this would mean that we would require 6 tape devices at certain moments.

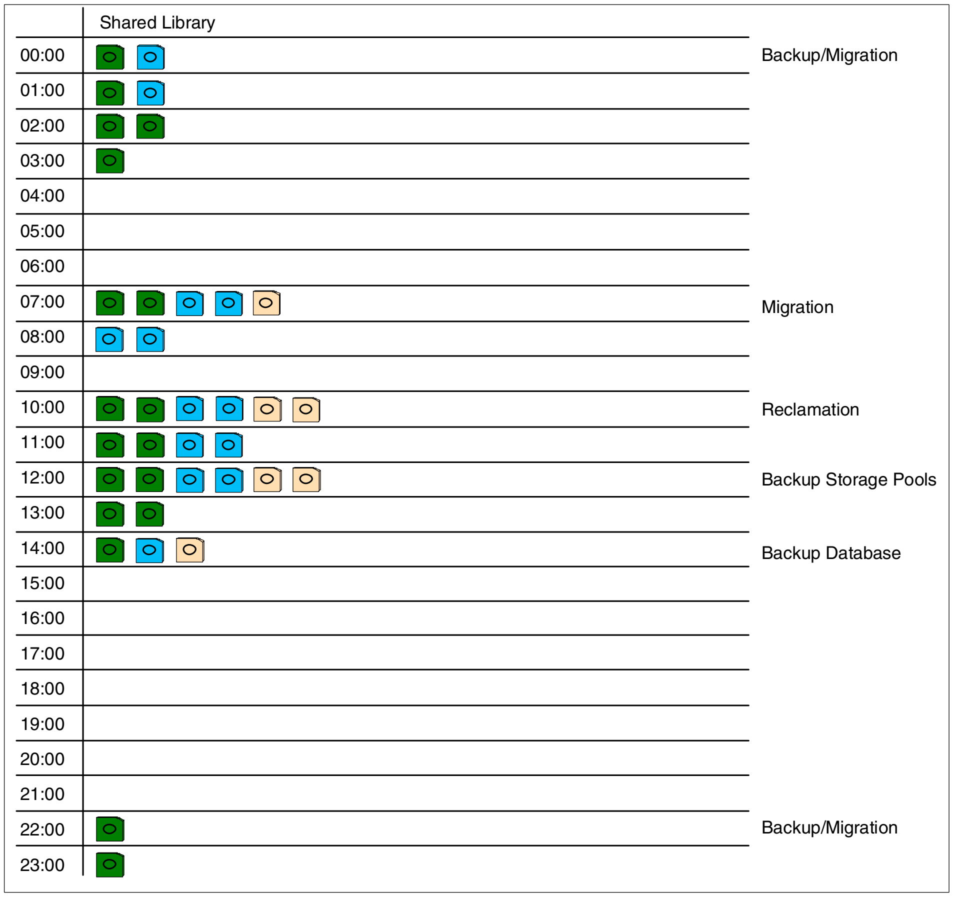

Figure 21. Total number of tape mounts

To improve the efficiency now, we could redistribute the usage of the tape devices over the total amount of time available. This can be done by rescheduling server processes. Figure 22 shows this rearrangement. As you can see, the utilization of the drives increases dramatically. Three drives would be sufficient to serve all processes.

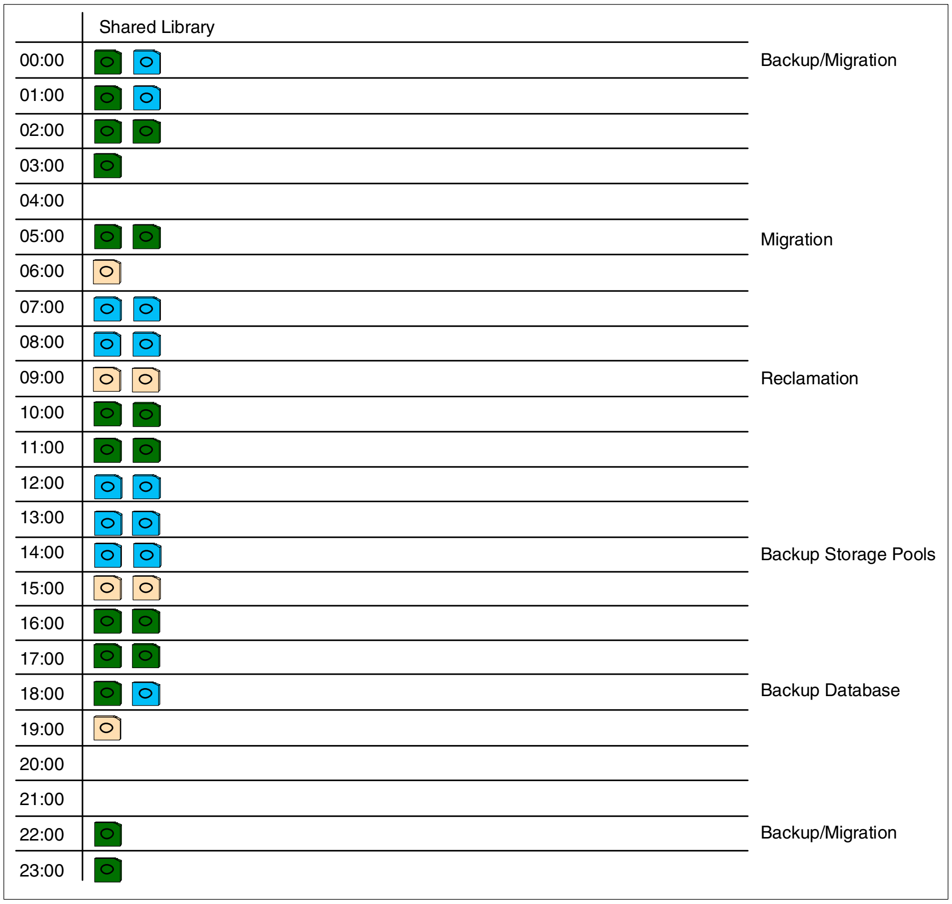

Figure 22. Redistributed tape mounts in the shared library

The rearrangement shown in Figure 22 would mean that all server processes on the different servers sharing the library must be precisely timed, and that a slight difference in duration of a process could lead to conditions where no more tape drives are available.

2.2.3.3 Conclusions

With the above in mind, one of the major things to keep in mind is the following: when sharing drives between servers, you are duplicating devices logically, but not physically. Do not simply reduce the number of drives to the minimum that you would need on one server. Have a number of drives that will be sufficient to serve all machines connected to the library.

Chapter 3. Building a SAN environment

Today, many sources give you a complete overview of various SAN configuration scenarios. In this chapter, we would like to discuss some possible scenarios for using SAN with Tivoli Storage Manager.

For detailed information about various SAN solutions, we recommend that you read the following IBM Redbooks:

- *Introduction to Storage Area Network, SAN,* SG14-5470
- *Tape future in fabrics*, SG24-5474

A SAN configuration can be customized to fit your specific needs. When building a SAN environment, we recommend that you always start with the most basic configuration possible, like connecting two or more devices together. Then you should verify the functionality of all your components. Once you have solved all of your problems, extend your configuration. We recommend this approach because SAN technology is still in its infancy; each component does not necessarily work with all other components.

.

Note:

The most recent information about SAN hardware, supported components and operating systems can be found on the following Web page:

`http://www.storage.ibm.com/ibmsan/products/sanfabric.htm`

Our project involved testing various Tivoli Storage Manager tape operations in our lab. For this purpose, we configured a basic SAN for each of the scenarios we tested.

This section details the steps we followed to configure each of the components that made up the SAN in our lab scenarios.

3.1 Host Bus Adapters (HBAs)

Host Bus Adapters are the means by which a particular system “talks” to the SAN. This section provides information on configuration of supported host bus adapters on their respective platforms.

3.1.1 RS/6000

The host bus adapter supported on the RS/6000 platform is categorized as Feature Code 6227, and is also known as the Emulex adapter. Although all the newer RS/6000 models are PCI-based machines, the Qlogic PCI adapters are **not** supported.

You should update your adapter code to the latest release, which is available from the following URL:

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/AIX/sdgaix.htm`

This Web page contains links to instructions for downloading, unpacking, and installing the adapter firmware. The StorWatch SAN Data Gateway Specialist for AIX is available here as well.

> **Important:**
>
> In order to make the 6227 adapter correctly identify SCSI devices attached to either the SAN Data Gateway or SAN Data Gateway Router, you *must* change the **max_xfer_size** attribute of the Fibre Channel adapter to **0x1000000**. Failure to do so will cause problems with identification and configuration of SCSI devices attached to the gateway or router. Use the procedure shown to change this attribute.

When initially installing or updating the firmware on a 6227 Fibre Channel adapter, it will be necessary to change the `max_xfer_size` attribute. Check the current setting using the `lsattr` command:

```
root@brazil:/ >lsattr -El fcs0
bus_intr_lvl  3          Bus interrupt level                               Fals
e
intr_priority 3          Interrupt priority                                Fals
e
bus_io_addr   0x7fffc00  Bus I/O address                                   Fals
e
bus_mem_addr  0x1ce9f000 Bus memory address                                Fals
e
max_xfer_size 0x100000   Maximum Transfer Size                             True
num_cmd_elems 200        Maximum number of COMMANDS to queue to the adapter True
pref_alpa     0x1        Preferred AL_PA                                   True
sw_fc_class   2          FC Class for Fabric                               True
init_link     al         INIT Link flags                                   True
root@brazil:/ >
```

Note that max_xfer_size is currently set to 0x100000, the default value. Update this value to 0x1000000 using the following procedure:

1. Shutdown the server, and disconnect the Fibre Channel connections to the card to be changed. This will keep the fscsiX devices from preventing the change.
2. Boot the system.
3. Change the max_xfer_size attribute using the chdev command (as in the following example).
4. Reconnect the Fibre Channel connections.
5. Run cfgmgr command, and verify that the tape devices are available:

```
root@brazil:/ >chdev -l fcs0 -a max_xfer_size=0x1000000
fcs0 updated
root@brazil:/ >
root@brazil:/ >lsattr -El fcs0
bus_intr_lvl  3          Bus interrupt level                                Fals
e
intr_priority 3          Interrupt priority                                 Fals
e
bus_io_addr   0x7fffc00  Bus I/O address                                    Fals
e
bus_mem_addr  0x1ce9f000 Bus memory address                                 Fals
e
max_xfer_size 0x1000000  Maximum Transfer Size                              True
num_cmd_elems 200        Maximum number of COMMANDS to queue to the adapter True
pref_alpa     0x1        Preferred AL_PA                                    True
sw_fc_class   2          FC Class for Fabric                                True
init_link     al         INIT Link flags                                    True
root@brazil:/ >
root@brazil:/ > cfgmgr
root@brazil:/ >
root@brazil:/ >lsdev -Cc tape
rmt0 Available 30-58-00-3,0 Other SCSI Tape Drive
rmt1 Available 10-68-01-2,0 IBM Magstar MP Tape Subsystem (FCP)
rmt3 Available 10-68-01-3,0 IBM Magstar MP Tape Subsystem (FCP)
root@brazil:/ >
```

At this point, the 6227 cards should see the gateway-attached SCSI devices.

3.1.2 Windows NT 4.0 Server and Windows 2000 Server

For Netfinity servers, the Qlogic 2100 and 2200 series Fibre Channel adapters are supported. Firmware and drivers for these adapters are available from the following URL:

http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/NT/sdgnt.htm

3.1.3 Sun Solaris

Sun Solaris servers support both the Qlogic QLA2100F HBA and the JNI FC64-1063 SBus adapter. Firmware and drivers for these adapters when attaching to a SAN Data Gateway can be found at the following URL:

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/Solaris/sdgsun.htm`

3.2 Fibre Channel interconnect devices

This section covers the setup of the actual Fibre Channel components.

SAN components are like building blocks; to a large extent, you can build just about anything you like, depending on what your needs are. Because of this, we will only cover the basic configuration of these devices; the information we present here is adequate to perform the test scenarios presented in Chapter 4, “Tape library sharing” on page 75. Refer to the User’s Guide for the device(s) you are using for more complete information.

3.2.1 Fibre Channel Hub 2107-H07 basic configuration

This product is an unmanaged 7-port hub which provides a single-point solution for Fibre Channel Arbitrated Loop (FC-AL) topologies and supports up to 100 MB per second data transmission speeds between servers and storage systems.

For more information about the 2107-H07 hub, visit the Web page:

`http://www.storage.ibm.com/hardsoft/products/fchub/fcshub.htm`

The 2103-H07 Fibre Channel hub is an entry level FC-AL hub with no configuration required for use. Preparation for this device consists of plugging Fibre Channel connections, and powering on.

3.2.2 SAN Data Gateway Router 2108-R03 basic configuration

The following URL contains the most recent support matrix for the SAN Data Gateway Router:

`http://www.storage.ibm.com/hardsoft/products/tape/ro3superserver.htm`

The router is currently *only* supported with Intel Windows based systems. Refer to the URL above for the latest support information.

The IBM SAN Data Gateway Router is very similar in appearance and function to the IBM SAN Data Gateway, with some functional restrictions. Refer to Chapter 1, “SAN overview” on page 1 for an overview of the differences between the SAN Data Gateway and the SAN Data Gateway Router products.

We began the setup of the gateway from a newly-initialized configuration. The purpose of this section is to provide an overview of the steps necessary to get a basic configuration working. For more information, or for more complex configurations, refer to the *Installation and User’s Guide 2108 Model R03*, SC26-7355-00.

1. Begin by readying either an ASCII terminal, or a terminal emulation program running on a PC, to talk to the service port (serial port) on the gateway. Set your communications for 19200 baud, 8 data bits, no parity, and xon/xoff flow control. A standard null modem DB-9 (female) to DB-9 (female) serial cable is necessary to communicate through the service port. Further details regarding service port communication cabling and settings are contained in the *Installation and User’s Guide 2108 Model R03*, SC26-7355.
2. When ready, apply power and turn on the gateway. The gateway should begin to display boot messages on the screen. The gateway uses a VxWorks real-time operating system, and provides a command line interface, as well as “out-band” network connectivity through the integrated ethernet port. When boot-up is complete, you should see a display similar to the following:

```
value = -1 = 0xffffffff
csSrvcInit
value = 0 = 0x0
amemInit
value = 0 = 0x0
scsintInit

SAN Router Version 0340.10 Built Mar 31 2000, 10:17:21
Router
0xc1f9dad0 (tShell): VPS Feature NOT Enabled
Data Mover is Disabled
SAN Management License is Valid
SCSI 2 - DE - Terminated
SCSI 1 - DE - Terminated
Router
IP Address set to 193.1.1.25:ffffff00
value = 0 = 0x0


Done executing startup script ffs0:sna.rc
- Service Port Reflected Signal Test
Starting shell
Router >
```

3. Once the gateway is booted, issue a SHOWBOX command to verify the hardware is as you expect it. Note that the router has only a single Fibre Channel interface. Note that the SHOWBOX output indicates FCOSW (Fibre Channel/Optical/Short Wave) in the Fibre Channel slot location.

```
Router > showBox
        _                                       _
 Front //|                                     //|
      // /------------------------------------// |
     // /                                    //  |
    // /                                    //   |
   // /                             /=====///    /
  ///====/                         /FCOSW///    /
 /========================================/    /
 | [DET=]                                 |   /
 |                               [FCOSW]  |  /
 | [DET=]                                 | /
 |________________________________________|/  Back

 DET = Differential SCSI with internal termination
 FCOSW = FibreChannel - Optical - Short Wave

 SCSI chan-1 requires Differential Cable
 SCSI chan-2 requires Differential Cable
 FibreChannel option card Slot 1 requires Short Wave Fiber Optic cable
```

4. Verify the firmware version of the router using the `VERSION` command. The latest firmware is available at:

   http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/sdgcdr.htm

   Instructions for updating the firmware are contained in the publication *Installation and User's Guide 2108 Model G07*, SC26-7304.

5. (optional) Set the name of the gateway with the `hostNameSet` command:

```
Router > hostNameSet "lab_SDGWR"
Write complete
Target hostname set to lab_SDGWR
value = 0 = 0x0
lab_SDGWR >
```

6. Set the IP address and network mask for the ethernet port using the `ethAddrSet` command:

```
lab_SDGWR > ethAddrSet "193.1.1.25", "255.255.255.0"
Inet Mask set to ffffff00 for Ethernet interface
Write complete
Host Address set to 193.1.1.25 for Ethernet interface
value = 0 = 0x0
lab_SDGWR >
```

7. Verify the address using the `ifShow` command:

```
lab_SDGWR > ifShow
lnPci (unit number 0):
     Flags: (0x63) UP BROADCAST ARP RUNNING
     Internet address: 193.1.1.25
     Broadcast address: 193.1.1.255
     Netmask 0xffffff00 Subnetmask 0xffffff00
     Ethernet address is 00:60:45:16:0f:35
     Metric is 0
     Maximum Transfer Unit size is 1500
     13 packets received; 0 packets sent
     0 input errors; 0 output errors
     0 collisions
lo (unit number 0):
     Flags: (0x69) UP LOOPBACK ARP RUNNING
     Internet address: 127.0.0.1
     Netmask 0xff000000 Subnetmask 0xff000000
     Metric is 0
     Maximum Transfer Unit size is 4096
     6 packets received; 6 packets sent
     0 input errors; 0 output errors
     0 collisions
value = 18 = 0x12
lab_SDGWR >
```

8. Verify that all attached SCSI devices can be seen by the gateway with the `targets` command. In the following screen listing, you can see the two 3570 tape drives as indexes 2 and 4, and the 3570 robotics are shown as index 3:

.

```
lab_SDGWR > targets
Idx Tdev       Vendor   Product          Rev  | Type Specific
----------------------------------------------|---------------------------------
--
  0 0xc0c60b00 PATHLGHT SAN Router       328a | Cmd/Cntrl Status 0h
  2 0xc1ffc590 IBM      03570C12         5424 | Tape: Blk Size 0 , flags 87h
  3 0xc1ffb290 IBM      03570C12         5424 | Changer: flags 7h
  4 0xc1f9e010 IBM      03570C12         5424 | Tape: Blk Size 0 , flags 87h
value = 4 = 0x4
lab_SDGWR >
```

9. Confirm your Fibre Channel connections using the `fcShow` command. In the following screen, you can see the single Fibre Channel interface:

```
lab_SDGWR > fcShow

----------------------------------------------------------------------
                       Fibre Channel Controllers
----------------------------------------------------------------------
 Ctlr : PCI Addr : ISP  :  Firmware   : Firmware : Loop :  Fabric  : Port
  Id  : Bs Dv Fn : Type :   State     : Version  :  ID  : Attached : Mode
----------------------------------------------------------------------
   1  : 00 06 00 : 2100 :   Ready     : 1.19.6   :   0  :    No    : Targ
----------------------------------------------------------------------

value = 76 = 0x4c = 'L'
lab_SDGWR >
```

10. In our lab configuration, each host was only attached to the gateway, that is, there were no other devices on our Fibre Channel loop. In the real world, this may very well not be the case. Fibre Channel, by default, uses a technique called *soft IDs* for addressing a given device on a loop.

    With soft IDs (also called *soft addressing*), the introduction of a new device, either host or target, could change the address of other devices on the loop. Some operating systems, like Windows, can deal relatively gracefully with an address change of this type. Others, including most UNIX variants, do not.

    In order to prevent address changes of this sort when a new device is introduced in a loop, you can (and should) use hard IDs, also called persistent IDs. These hard IDs can be assigned using the `setFcHardId` command. In the following example, we have set hard IDs for our Fibre Channel interface:

```
lab_SDGWR > setFcHardId 1,0
value = 0 = 0x0
lab_SDGWR >
```

11. (optional) In order to prevent Windows machines from “seeing” the SAN Data Gateway as a device, and attempting to locate an appropriate driver for it (there is none needed), you need to “hide” the Command and Control interface from the Windows operating system. Use the `disableCC` command to do this.

    If you do not do this, the gateway will still function fine. However, Windows will see the Command and Control interface on LUN 0, and prompt you to install drivers. Since there are no drivers, Windows will show the SAN Data Gateway with a question mark (**?**) next to it in Device Manager. Running `disableCC 2` hides LUN 0, and prevents this prompting:

```
lab_SDGWR > disableCC 2
value = 1 = 0x1
lab_SDGWR >
```

12. Define the host types attached to the Fibre Channel interfaces on the gateway. You can display the current settings using the `hostTypeShow` command.

    Host types that valid for the gateway are:

    - NT
    - AIX
    - Solaris
    - HP-UX
    - Netware
    - Gateway
    - Switch
    - Generic

```
lab_SDGWR > hostTypeShow
FC 1: Type 1 - nt
value = 0 = 0x0
lab_SDGWR >

lab_SDGWR > setHost 1, "switch"
value = 0 = 0x0
lab_SDGWR >
```

13. (optional) Add attached hosts to the gateway's internal hosts table using the `host "add"` command. The following listing shows the syntax for adding hosts to the table; repeat for each host to be added. The contents of the table can be verified using the `host "list"` command:

```
lab_SDGWR > host "add","brazil","193.1.1.11"
Host file missing. Creating host file.
value = 0 = 0x0
lab_SDGWR > host "add","jamaica","193.1.1.13"
value = 0 = 0x0
lab_SDGWR > host "add","diomede","193.1.1.16"
value = 0 = 0x0
lab_SDGWR >
lab_SDGWR > host "list"
193.1.1.11 brazil
193.1.1.13 jamaica
193.1.1.16 diomede
value = 0 = 0x0
lab_SDGWR >
```

14.After setting the parameters listed above, reboot the gateway using the `reboot` command:

```
Lab_SDGW1 > reboot
==  reboot
c00dc048 - _reboot
c011f41c - _yystart
c0120f54 - _yyparse
c00dad0c - _execute
c00daba8 - _shell
c00da9a8 - _shell
c011876c - _vxTaskEntry
VCM: 20 tasks started
MGT List Count 0
CMD List Count 0
LBuffer Pools:
 2097152 bytes: Free   8 of   8 - Wait   0,  0, Sig   0 N   0 L   0 list 0xc01e
  262144 bytes: Free 128 of 128 - Wait   0,  0, Sig   0 N   0 L   0 list 0xc01e
   32768 bytes: Free 256 of 256 - Wait   0,  0, Sig   0 N 227 L 227 list 0xc01e
    4096 bytes: Free 648 of 1664 - Wait   0,  0, Sig   0 N 1038 L  22 list 0xc0

Syncing volume ffs0. Volume now unmounted.
Flash File Systems now synced.
```

This completes a basic configuration of the SAN Data Gateway Router.

The SAN Data Gateway Router can also be managed using the StorWatch SAN Data Gateway Specialist; refer to Section A.1, “StorWatch SAN Data Gateway Specialist” on page 261 for instructions on installing and configuring this tool.

3.2.3 SAN Data Gateway 2108-G07 basic configuration

The IBM Storage Area Network (SAN) Data Gateway is a hardware solution that enables the attachment of SCSI storage systems to Fibre Channel adapters on specific Intel-based servers running Windows NT and UNIX-based servers from IBM and Sun.

Requirements for the SAN Data Gateway listed by operating system are available at the following URL:

`http://www.storage.ibm.com/hardsoft/products/tape/tapesupport.htm`

When we prepared the SAN Data Gateway (SDGW) for scenario 2, we began with a “clean” initialized box, the same as what you would receive if you had purchased a new unit.

The purpose of this section is to provide an overview of what is necessary to get a basic configuration working; for more detailed information, and information on more complex configurations, refer to *Installation and User’s Guide 2108 Model G07*, SC26-7304.

1. Begin by readying either an ASCII terminal, or a terminal emulation program running on a PC to talk to the service port (serial port) on the gateway. Set your communications for 19200 baud, 8 data bits, no parity, and xon/xoff flow control. A standard null modem DB-9 (female) to DB-9 (female) serial cable is necessary to communicate through the service port. Further details regarding service port communication cabling and settings are contained in the *Installation and User’s Guide 2108 Model G07*, SC26-7304.
2. When ready, apply power and turn on the gateway. The gateway should begin to display boot messages on the screen. The gateway uses a VxWorks real-time operating system, and provides a command line interface, as well as “out-band” network connectivity through the integrated ethernet port. When bootup is complete, you should see a display similar to the following:

```
SAN Gateway Version 0340.16 Built May 19 2000, 08:27:01
Gateway
0xc1f9e2b0 (tShell): VPS Feature NOT Enabled
Data Mover is Disabled
SCSI 3 - DE - Terminated
SCSI 1 - DE - Terminated
SCSI 4 - DE - Terminated
SCSI 2 - DE - Terminated
sdgw1
IP Address set to 193.1.1.25:ffffff00
value = 0 = 0x0


Done executing startup script ffs0:sna.rc
- Service Port Reflected Signal Test
Starting shell
Gateway >
```

3. Once the gateway is booted, issue a `showBox` command to verify the hardware is as you expect it. Note that our gateway had two dual-port Fibre Channel interfaces (PMCs) installed. In a gateway with single-port PMCs, the diagram returned from a `showBox` command will indicate FCOSW instead of DFCSW in the Fibre Channel slot locations.

```
sdgw1 > showBox
        _                                       _
 Front //|     /-------------------------------//|
      // /----/          DDF Board           // |
     // /    /                              //  |
    // /     -------------------------------//   |
   // /                  /=====/ /=====///    /
  ///====/ /====/          /DFCSW/ /DFCSW///    /
 /=========================================/    /
 | [DET=] [DET=]                           |   /
 |               [=3===] [DFCSW] [DFCSW]   |  /
 | [DET=] [DET=]                           | /
 |_________________________________________|/  Back

 DET = Differential SCSI with internal termination
 DFCSW = Dual FibreChannel - Optical - Short Wave


 SCSI chan-1 requires Differential Cable
 SCSI chan-2 requires Differential Cable
 SCSI chan-3 requires Differential Cable
 SCSI chan-4 requires Differential Cable
 FibreChannel option card Slot 1 Ports 1,4 require Short Wave Fiber Optic cable
 FibreChannel option card Slot 2 Ports 2,5 require Short Wave Fiber Optic cable
```

4. Verify the firmware level of the gateway, using the `version` command. The version of firmware running on this gateway is 3.40.16, which is the latest version available as of this writing. The latest firmware can be obtained from the following URL:

   `http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/sdgcdr.htm`

   For instructions on updating the firmware, refer to *Installation and User's Guide 2108 Model G07*, SC26-7304.

```
Lab_SDGW1 > version

SAN Gateway Version 0340.16 Built May 19 2000, 08:27:01
VxWorks (for Pathlight (i960RD)) version 5.3.1.
Kernel: WIND version 2.5.
value = 26 = 0x1a
```

5. (optional) Set the name of the gateway with the `hostNameSet` command:

```
Gateway > hostNameSet "Lab_SDGW1"
Write complete
Target hostname set to Lab_SDGW1
value = 0 = 0x0
Lab_SDGW1 >
```

6. Set the IP address and network mask for the ethernet port using the `ethAddrSet` command:

```
Lab_SDGW1 > ethAddrSet "193.1.1.25","255.255.255.0"
Inet Mask set to ffffff00 for Ethernet interface
Write complete
Host Address set to 193.1.1.25 for Ethernet interface
value = 0 = 0x0
```

7. Verify the address using the `ifShow` command:

```
Lab_SDGW1 > ifShow
lnPci (unit number 0):
     Flags: (0x63) UP BROADCAST ARP RUNNING
     Internet address: 193.1.1.25
     Broadcast address: 193.1.1.255
     Netmask 0xffffff00 Subnetmask 0xffffff00
     Ethernet address is 00:60:45:16:0b:a3
     Metric is 0
     Maximum Transfer Unit size is 1500
     24 packets received; 2 packets sent
     0 input errors; 0 output errors
     0 collisions
lo (unit number 0):
     Flags: (0x69) UP LOOPBACK ARP RUNNING
     Internet address: 127.0.0.1
     Netmask 0xff000000 Subnetmask 0xff000000
     Metric is 0
     Maximum Transfer Unit size is 4096
     6 packets received; 6 packets sent
     0 input errors; 0 output errors
     0 collisions
value = 18 = 0x12
```

8. Verify that all attached SCSI devices can be seen by the gateway with the `targets` command. In the following screen listing, you can see the two 3570 tape drives as indexes 2 and 4, and the 3570 robotics are seen as index 3.

```
Lab_SDGW1 > targets
Idx Tdev       Vendor   Product          Rev  | Type Specific
--------------------------------------------|--------------------------------
  0 0xc1745b80 PATHLGHT SAN Gateway      0016 | Cmd/Cntrl Status 0h
  2 0xc13a9f10 IBM      03570C12         5424 | Tape: Blk Size 0 , flags 87h
  3 0xc13a7b90 IBM      03570C12         5424 | Changer: flags 7h
  4 0xc1ffc390 IBM      03570C12         5424 | Tape: Blk Size 0 , flags 87h
value = 4 = 0x4
```

9. Confirm your Fibre Channel connections using the `fcShow` command. The following screen shows that the three attached hosts are seen by the gateway.

```
Lab_SDGW1 > fcShow


-----------------------------------------------------------------------
                        Fibre Channel Controllers
-----------------------------------------------------------------------
 Ctlr : PCI Addr : ISP  :  Firmware   : Firmware : Loop :  Fabric  : Port
  Id  : Bs Dv Fn : Type :   State     : Version  :  ID  : Attached : Mode
-----------------------------------------------------------------------
   1  : 00 06 00 : 2200 :   Ready     : 2.01.12  :   0  :    No    : Targ
   2  : 00 07 00 : 2200 :   Ready     : 2.01.12  :   0  :    No    : Targ
   4  : 00 18 00 : 2200 :   Ready     : 2.01.12  :   0  :    No    : Targ
-----------------------------------------------------------------------

value = 76 = 0x4c = 'L'
```

10.In our lab configuration, each host was only attached to the gateway, that is, there were no other devices on our Fibre Channel loops. In the real world, this may very well not be the case. Fibre channel, by default, uses a technique called *soft IDs* for addressing a given device on a loop.

With soft IDs (also called *soft addressing*), the introduction of a new device, either host or target, could change the address of other devices on the loop. Some operating systems, like Windows, can deal relatively gracefully with an address change of this type. Others, including most UNIX variants, do not.

In order to prevent address changes of this sort when a new device is introduced in a loop, you can (and should) use hard IDs, also called persistent IDs. These hard IDs can be assigned using the `setFcHardId` command. In the following example, we have set hard IDs for all of our Fibre Channel interfaces:

```
Lab_SDGW1 > setFcHardId 1,0
value = 0 = 0x0
Lab_SDGW1 > setFcHardId 2,0
value = 0 = 0x0
Lab_SDGW1 > setFcHardId 4,0
value = 0 = 0x0
Lab_SDGW1 > setFcHardId 5,0
value = 0 = 0x0
```

11.(optional) In order to prevent Windows machines from “seeing” the SAN Data Gateway as a device, and attempting to locate an appropriate driver for it (there is none needed), you need to “hide” the Command and Control interface from the Windows operating system. Use the `disableCC` command to do this.

If you do *not* do this, the gateway will still function fine. However, Windows will see the Command and Control interface on LUN 0, and will prompt you to install drivers. Since there are no drivers, Windows will show the SAN Data Gateway with a question mark (**?**) next to it in Device Manager. Running `disableCC 2` hides LUN 0, and prevents this prompting.

```
Lab_SDGW1 > disableCC 2
value = 1 = 0x1
```

12. Define the host types attached to the Fibre Channel interfaces on the gateway. You can display the current settings using the `hostTypeShow` command.

    Host types that valid for the gateway are:

    - NT
    - AIX
    - Solaris
    - HP-UX
    - Netware
    - Gateway
    - Switch
    - Generic

    The default type for each port is NT. Change the host type setting for each port as required, using the `setHost` command, as follows:

```
Lab_SDGW1 > hostTypeShow
FC 1: Type 1 - nt
FC 2: Type 1 - nt
FC 4: Type 1 - nt
FC 5: Type 1 - nt
value = 0 = 0x0
Lab_SDGW1 >
Lab_SDGW1 > setHost 4,"AIX"
value = 0 = 0x0
Lab_SDGW1 > hostTypeShow
FC 1: Type 1 - nt
FC 2: Type 1 - nt
FC 4: Type 2 - aix
FC 5: Type 1 - nt
value = 0 = 0x0
Lab_SDGW1 >
```

13. After setting the parameters listed above, reboot the gateway using the `reboot` command:

```
Lab_SDGW1 > reboot
==  reboot
c00dc048 - _reboot
c011f41c - _yystart
c0120f54 - _yyparse
c00dad0c - _execute
c00daba8 - _shell
c00da9a8 - _shell
c011876c - _vxTaskEntry
VCM: 20 tasks started
MGT List Count 0
CMD List Count 0
LBuffer Pools:
 2097152 bytes: Free   8 of   8 - Wait   0,  0, Sig   0 N   0 L   0 list 0xc01e
  262144 bytes: Free 128 of 128 - Wait   0,  0, Sig   0 N   0 L   0 list 0xc01e
   32768 bytes: Free 256 of 256 - Wait   0,  0, Sig   0 N 227 L 227 list 0xc01e
    4096 bytes: Free 648 of 1664 - Wait   0,  0, Sig   0 N 1038 L  22 list 0xc0

Syncing volume ffs0. Volume now unmounted.
Flash File Systems now synced.
```

14. (optional) Add attached hosts to the gateway's internal hosts table using the `host "add"` command. The following listing shows the syntax for adding hosts to the table; repeat for each host to be added. The contents of the table can be verified using the `host "list"` command:

```
Lab_SDGW1 > host "add","brazil","193.1.1.11"
Host file missing. Creating host file.
value = 0 = 0x0
Lab_SDGW1 > host "add","jamaica","193.1.1.13"
value = 0 = 0x0
Lab_SDGW1 > host "add","diomede","193.1.1.16"
value = 0 = 0x0
Lab_SDGW1 > host "list"
193.1.1.11 brazil
193.1.1.13 jamaica
193.1.1.16 diomede
value = 0 = 0x0
```

This completes the basic configuration steps necessary for the 2108-G07 SAN Data Gateway.

The SAN Data Gateway can also be managed using the StorWatch SAN Data Gateway Specialist; refer to Section A.1, “StorWatch SAN Data Gateway Specialist” on page 261 for instructions on installing and configuring this tool.

3.2.4 IBM Fibre Channel Switch 2109-S08 basic configuration

The most recent list of requirements for the SAN Fibre Channel Switch, listed by host type, can be found at:

`http://www.storage.ibm.com/hardsoft/products/fcswitch/supserver.htm`

Firmware for the SAN Fibre Channel Switch is available at:

`http://www.storage.ibm.com/hardsoft/products/fcswitch/download.htm`

We began with a “clean” switch configuration, that is no prior configurations stored on the switch.

1. Begin by readying either an ASCII terminal, or a terminal emulation program running on a PC to talk to the service port (serial port) on the switch. Set your communications for 19200 baud, 8 data bits, no parity, and xon/xoff flow control. A standard straight-through DB-9 (female) to DB-9 (female) serial cable is necessary to communicate through the service port. Further details regarding service port communication cabling and settings are contained in the *IBM SAN Fibre Channel Switch 2109 Model S08 Users Guide,* SC26-7349.
2. When ready, apply power and turn on the switch. The gateway should begin to display boot messages on the screen. The gateway uses a VxWorks real-time operating system, and provides a command line interface, as well as “out-band” network connectivity through the integrated ethernet port. When bootup is complete, you should see a display similar to the following:

```
Fabric OS (tm)  Release v2.1.7

WARNING: switch still initializing - please wait

switch1:admin>
Running System DRAM Test ...... passed.

Running Port Register Test .... passed.

Running Central Memory Test ... passed.

Running CMI Test ............. passed.

Running CAM Test ............. passed.

Running Port Loopback Test .... passed.
Sep 26 14:19:02.433  port 1: FLOGI before switch ready, disabling port
10  9  8  7  6  5  4  3  2  1
fabric: Principal switch
fabric: Domain 2
http: server initialized
ns_init: Name Server initialized
as_init: Alias Server initialized
sesd enabled
Sep 26 14:19:16.083  port 1: FLOGI 0x021100 cos=0x8 bb_credit=2 df_size=2048 cf=0x0
Sep 26 14:19:16.099  port 1: PLOGI s_id=0x21100 d_id=0xfffffc cos=0x8 df_size=2048
Sep 26 14:19:17.083  port 1: PLOGI s_id=0x21100 d_id=0xfffc02 cos=0x8 df_size=2048
Sep 26 14:19:18.149  port 3: FLOGI 0x0213ef cos=0x8 bb_credit=0 df_size=1024 cf=0x800
Sep 26 14:19:18.166  port 6: FLOGI 0x0216ef cos=0x8 bb_credit=0 df_size=2048 cf=0x800
Sep 26 14:19:18.183  port 6: PLOGI s_id=0x216ef d_id=0xfffffc cos=0x8 df_size=2048
Sep 26 14:19:18.183  port 3: PLOGI s_id=0x213ef d_id=0xfffffc cos=0x8 df_size=1024

switch1:admin>
```

3. Once the switch is booted, and you get a command prompt, set the IP parameters for the network interface using the `IpAddrSet` command:

```
switch1:admin> ipAddrSet
Ethernet IP Address [193.1.1.14]:
Ethernet Subnetmask [255.255.255.0]:
Fibre Channel IP Address [none]:
Fibre Channel Subnetmask [none]:
Gateway Address [9.113.25.254]:
switch1:admin>
```

4. Run the `switchShow` command to find out what the switch can see:

```
switch1:admin> switchShow
switchName:     switch1
switchType:     3.4
switchState:    Online
switchRole:     Principal
switchDomain:   2
switchId:       fffc02
switchWwn:      10:00:00:60:69:20:1d:4e
port  0: sw  No_Light
port  1: sw  Online     F-Port  21:00:00:e0:8b:02:36:29
port  2: sw  No_Light
port  3: sw  Online     L-Port  1 public
port  4: id  No_Light
port  5: sw  No_Light
port  6: id  Online     L-Port  1 public
port  7: id  No_Light
switch1:admin>
```

   In the above screen shot, we have three ports showing as “online”.

5. (optional) Set the name of the switch to something descriptive, using the `switchName` command:

```
switcha:admin> switchName "switch1"
Updating flash ...
switch1:admin>
switch1:admin> switchName
switch1
switch1:admin>
```

The 2109-series switches auto sense the port type and host operating system, so unlike the SAN Data Gateway and router products, it is not necessary to set the host type for each port.

At this point, the switch is set up for non-zoned operation.

While command line configuration and operation of the switch is certainly possible, IBM provides the IBM StorWatch Fibre Channel Switch Specialist free of charge, which is a far easier method of configuring and managing the switch. An overview of the setup and use of this product is provided in Section A.2, “StorWatch Fibre Channel Switch Specialist” on page 269. Further information is available in the publication *IBM SAN Fibre Channel Switch 2109 Model S08 Users Guide,* SC26-7349.

3.2.5 Vicom SLIC Router FC-SL basic configuration

The most recent information about supported SAN devices can be found at:

http://www.vicom.com/product_fcsl.html

Begin the preparation of the SLIC Router with clearing the Node Mapping Table:

1. Power off the SLIC Router.
2. Interconnect the SSA ports A1 and A2 with an SSA cable.
3. Power on the SLIC Router.
4. Wait for the completion code “060” which is indicated by the LEDs on the rear panel of the SLIC Router.

   Refer to the SLIC Router documentation for further details on how the operation codes are presented through the LEDs.
5. Run the Subsystem Diagnostic Test:

- Power down the SLIC Router, set all the switches in SW1 to the down position, set the switches 7-2 to the upper position, and set switches 1,0 to the down position (this is called mode 3).
- Attach the SSA loop to the SLIC Router. Note that if you want to connect more than one drawer, you have to daisy chain these drawers together.
- Now you can power on the SLIC Router. You can see that the green status LED blinks fast. The SLIC Router now tries to identify all the SSA drives. Once this action is completed, the status LED will stop blinking fast.
- Power off the SLIC Router. Set the switches in SW2 following: switches 7-4 to the upper position, switches 3-0 to the down position. Then power on the SLIC Router.

6. Create physical drives:

- Power off the SLIC Router. Set switch SW1 to assign unique ID. In our environment we assign the ID 5 to our SLIC Router. To assign this value, set the dip switches 0 and 2 to the down position.

  See the SLIC Router Installation and User Guide to find out what ID number is represented by each dip switch.
- In the switch SW2, set the switches 7-2 to the upper position, switches 1 and 0 to the down position.
- Power up the SLIC Router.

Now all the drives on the SSA loop are defined as mapped physical drives, which means you can access these drives from the host system.

For more detailed information about different configurations with the SLIC Router, read the documentation supplied with this device.

Also see Appendix B, “Configuring SLIC Manager on Windows” on page 273.

Chapter 4. Tape library sharing

This chapter describes the tape library sharing functionality provided in Tivoli Storage Manager 3.7 and later releases, and how to implement this feature in a typical environment. In particular, we cover two different methods for sharing a library using Tivoli Storage Manager:

- SCSI tape library sharing
- 3494 tape library sharing

Introduction to SCSI tape library sharing

In releases of Tivoli Storage Manager prior to version 3.7, sharing the resources of a SCSI attached tape library was possible only through the use of server-server virtual volumes. Using this method, only one server is physically attached to the library; other servers transferred data to the library-attached server via a LAN connection. This allowed other Tivoli Storage Manager servers to store and access data on a given library, but had the drawback that it required that data stored on the library be passed over the network connection twice: once when the data was passed from client to server, and again when passed from source-server to target-server.

Starting with Tivoli Storage Manager Version 3.7, SCSI tape library sharing with Tivoli Storage Manager involves using a direct SAN connection from any participating server to the tape library to handle the data transfer from multiple Tivoli Storage Manager servers to a shared library device. Figure 23 shows a graphical representation of a typical shared SCSI library scenario.

In this kind of shared environment, one of the Tivoli Storage Manager servers acts as a *library manager*, owning and controlling the tape library device. All the other Tivoli Storage Manager servers sharing this library act as *library clients*. A library client requests shared library resources as drives or media from the Tivoli Storage Manager server acting as library manager, but uses the resources independently. The library manager coordinates the access to these resources.

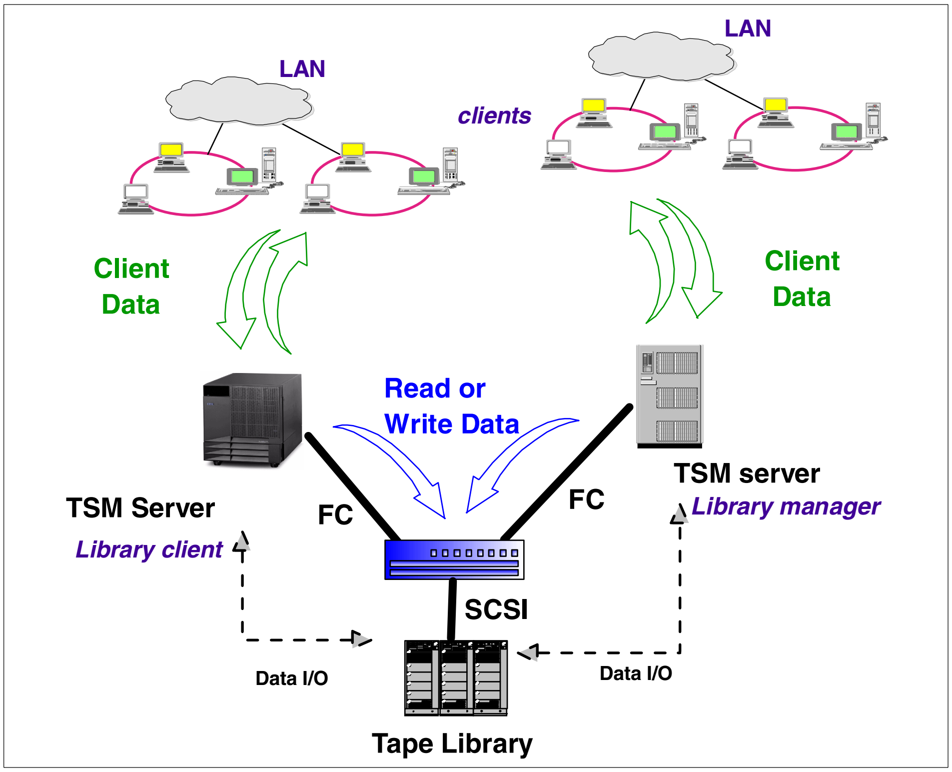

Figure 23. Overview of SCSI library sharing

Sharing a SCSI library between two Tivoli Storage Manager servers using the the library sharing functionality provided in 3.7 and higher has the following advantages:

- This eliminates the necessity to pass client data from one server to another over the LAN for server-server virtual volumes. By using a direct SAN connection, data to be stored in the library is written directly to the library drives via the SAN.
- There is a more efficient use of scratch volumes, as they can be shared between different Tivoli Storage Manager servers attached to the library.
- System resources remain free on the target server, as the source server does not need to pass the data through the target server to reach the library.

SCSI library sharing is addressed in Section 4.1, “SCSI tape library sharing” on page 77.

Introduction to 3494 library sharing

While sharing of SCSI libraries is new, the ability to share the resources of a 3494 library has existed for some time. Usage of a 3494 library's resources in the fashion described in the previous section is actually referred to as "partitioning" in the 3494 environment. Refer to Section 4.2.2, "Terminology: partitioning versus sharing" on page 112 for a discussion of the differences between "sharing" and "partitioning" in the 3494 environment.

Due to the enhanced capabilities of IBM's 3494 library introduced by its built-in library manager as compared to traditional SCSI libraries, there is no need to define a library manager in Tivoli Storage Manager. Access to library resources is managed by the IBM 3494 library manager itself.

New for Release 3.7 and higher is support for twin-tailed 3590 SCSI drives in the 3494. This allows multiple Tivoli Storage Manager servers to be connected to individual SCSI 3590 drives. Prior to the addition of this feature, called *auto-sharing*, individual drives in a 3494 library could only be attached to a particular server. The individual drives were "owned" only by that server.

The new 3590 Fibre Channel attached drives are also supported in Release 4.1.0.

Because of these differences, configuration and operation of the 3494 is treated as a separate topic in Section 4.2, "Sharing an IBM 3494 library" on page 110.

General requirements

The requirements for library sharing with Tivoli Storage Manager are listed at the following URL:

`http://www.tivoli.com/support/storage_mgr/requirements.html`

4.1 SCSI tape library sharing

This section explains the terminology used in a shared SCSI environment, and demonstrates the steps necessary to implement a shared SCSI library under different operating system platforms.

The following paragraphs give a short definition of the library manager and the library client.

Library manager

Tape library devices (robotics, drives) can only be accessed by one server at a time. Due to this fact, it is necessary to have one central “library manager” or server that controls who has control of a given library device. The library manager is the Tivoli Storage Manager server that determines which server has access to a library device at a given time.

The library manager also keeps track of the ownership, status and location of cartridges or volumes in the library. Library inventory information is maintained on the library manager server; library clients request access to volumes through the library manager.

Library client

The library client, which is simply another Tivoli Storage Manager server attached to the library, must send a request to the primary library manager when it requires access to library facilities. It does not have direct access to the library robotics in the same fashion as the library manager does.

When the library client requires access to a tape, whether the tape is scratch or private, it sends a request to the library manager, which mounts the tape, and updates ownership of the volumes as needed.

If the library client performs certain operations, such as an “audit library”, no tape mounts are actually performed. Instead, the library client’s database is checked against the library manager’s database.

Information regarding operations with a shared SCSI library can be found in Section 4.1.7, “Using a shared SCSI library with Tivoli Storage Manager” on page 103.

4.1.1 Steps for SCSI tape library sharing — an overview

This section covers the setup of a SAN-attached shared SCSI library under different operating systems. We will present three different scenarios, and the steps necessary to configure each:

1. A 3570 SCSI library shared using an unmanaged model 2103-H07 Hub, in conjunction with a SAN Data Gateway Router, model 2108-R03:
    a. Prepare SCSI library (Section 4.1.2.1, “SCSI library preparation” on page 80).
    b. Configure Fibre Channel Hub (Section 4.1.3.1, “IBM Fibre Channel Hub 2103-H07 preparation” on page 95).
    c. Configure SAN Data Gateway Router (Section 4.1.3.2, “IBM SAN Data Gateway Router 2108-R03 preparation” on page 95).

   d. Prepare server hardware (Section 3.1, “Host Bus Adapters (HBAs)” on page 51), including HBA microcode and device driver update/installation.

   e. Prepare server operating system and load drivers (Section 4.2.3.4, “Operating system preparation” on page 114).

   f. Perform Tivoli Storage Manager configuration (Section 4.2.4, “Tivoli Storage Manager configuration” on page 118).

2. A 3570 SCSI library shared between servers using an IBM SAN Data Gateway, model 2108-G07:

   a. Prepare SCSI library (Section 4.1.2.1, “SCSI library preparation” on page 80).

   b. Configure SAN Data Gateway (Section 4.1.2.2, “System and SAN hardware preparation” on page 80).

   c. Prepare server hardware (Section 3.1, “Host Bus Adapters (HBAs)” on page 51), including HBA microcode and device driver update/installation.

   d. Prepare server operating system and load drivers (Section 4.2.3.4, “Operating system preparation” on page 114).

   e. Perform Tivoli Storage Manager configuration (Section 4.2.4, “Tivoli Storage Manager configuration” on page 118).

3. A 3570 SCSI library shared via a Fibre Channel switch, model 2109-S08, and SAN Data Gateway, model 2108-G07:

   a. Prepare SCSI library (Section 4.1.2.1, “SCSI library preparation” on page 80).

   b. Configure SAN Fibre Channel switch (Section 4.1.2.2, “System and SAN hardware preparation” on page 80).

   c. Prepare server hardware (Section 3.1, “Host Bus Adapters (HBAs)” on page 51), including HBA microcode and device driver update/installation.

   d. Prepare server operating system and load drivers (Section 4.2.3.4, “Operating system preparation” on page 114).

   e. Perform Tivoli Storage Manager configuration (Section 4.2.4, “Tivoli Storage Manager configuration” on page 118).

Note:

Much of the information contained in this section pertains only to SCSI tape libraries. For information on sharing for 3494 libraries, see Section 4.2, “Sharing an IBM 3494 library” on page 110

4.1.2 Common steps for library scenarios

This section covers steps common to all three shared SCSI library scenarios.

4.1.2.1 SCSI library preparation

Preparation of a SCSI attached library which is to be used in a shared environment is identical to what you would do in a non-shared environment.

In the case of the 3570 Magstar library we used, no changes were necessary. We simply left the library in base and random modes, the same as would be used for direct SCSI attachment to a single Tivoli Storage Manager server.

Refer to the installation and user guides for the particular make and model of the library that you are using for the necessary setup instructions.

4.1.2.2 System and SAN hardware preparation

Information pertaining to hardware configuration and firmware levels relevant to the illustrated scenarios are contained in Section 3.1, “Host Bus Adapters (HBAs)” on page 51, and in Section 3.2, “Fibre Channel interconnect devices” on page 54.

4.1.2.3 Operating system preparation

There are certain preparatory steps for the different operating systems we used that are common to the three scenarios covered. This section details these steps.

AIX preparation

For an RS/6000 server implementation, AIX version 4.3.3.0.04 is required for library sharing.

On an RS/6000 server, the supported adapter for Fibre Channel attachment is the adapter with FC 6227. You should verify that the Fibre Channel adapters in your machine are at the latest available firmware level. The latest adapter microcode, as well as information on finding your current microcode level, can be found at the following URL:

`http://www.rs6000.ibm.com/support/micro/download.html`

Additionally, the AIX Atape device driver is also needed, which is available from the following URL:

`ftp://index.storsys.ibm.com/devdrvr/AIX`

As of this writing, the most recent release for Atape is Release 5.3.9.0.

After a fresh install of AIX 4.3.3, we installed the latest available maintenance level, bring the machine up to 4.3.3.0.04. After completion, the following levels are displayed:

```
root@brazil:/ >lslpp -l all | grep FC
                            4.3.3.10  COMMITTED  Common IBM FC Software
  devices.fcp.tape.rte      4.3.3.0   COMMITTED  FC SCSI Tape Device Software
  devices.pci.df1000f7.com  4.3.3.15  COMMITTED  Common PCI FC Adapter Device
                            4.3.3.11  COMMITTED  PCI FC Adapter Device
  devices.pci.df1000f7.rte  4.3.3.0   COMMITTED  PCI FC Adapter Device Software
  devices.fcp.tape.rte      4.3.3.0   COMMITTED  FC SCSI Tape Device Software
  devices.pci.df1000f7.com  4.3.3.0   COMMITTED  Common PCI FC Adapter Device
  devices.pci.df1000f7.rte  4.3.3.0   COMMITTED  PCI FC Adapter Device Software
```

Now:

1. Shut down the server, and turn off both the SCSI library and the gateway
2. Connect the SCSI cables to the SAN Data Gateway
3. Connect the Fibre Channel cables from the server to the gateway
4. Turn on the SCSI library, and wait for self test to complete
5. Turn on the SAN Data Gateway, and wait for it to finish booting
6. Boot the AIX server

At this point, the drives in the SCSI library should be visible.

```
root@brazil:/ >lsdev -Cc tape
rmt0 Available 10-68-01 Other FC SCSI Tape Drive
rmt1 Available 10-68-01 Other FC SCSI Tape Drive
```

Note, however, that while the system is able to “see” the drives, they are not yet identified correctly. It is still necessary to install the “Atape” driver.

Download the Atape driver to an temporary directory, and install using the `smitty` command. After completion, the original tape devices will need to be removed via the `rmdev` command, and added again with `cfgmgr` command:

```
root@brazil:/ >rmdev -l rmt0 -d
rmt0 deleted
root@brazil:/ >rmdev -l rmt1 -d
rmt1 deleted
root@brazil:/ >lsdev -Cc tape
root@brazil:/ >
root@brazil:/ >cfgmgr
root@brazil:/ >
root@brazil:/ >lsdev -Cc tape
rmt0 Available 10-68-01 IBM Magstar MP Tape Subsystem (FCP)
rmt2 Available 10-68-01 IBM Magstar MP Tape Subsystem (FCP)
root@brazil:/ >
```

Notice that the tape drives are now correctly identified as Magstar drives. This confirms that the Atape driver is loaded correctly.

At this point, you should proceed to install Tivoli Storage Manager. We recommend that you install the base 4.1.0 code from CD, the upgrade to the required 4.1.1 code, which is available from:

`ftp://index.storsys.ibm.com/tivoli-storage-management/maintenance/server/v4r1/`

When you install Tivoli Storage Manager, be sure to install Fibre Channel support by selecting the `tivoli.tsm.devices.fcprte` file set, as this is required for support of Fibre Channel connected devices.

Note:

The fileset `tivoli.tsm.devices.rte` is a subset of the `tivoli.tsm.devices.fcprte` fileset. Installation of the `tivoli.tsm.devices.fcprte` fileset provides both SCSI and Fibre Channel device support. During installation, you must select one or the other; selecting both will cause the installation of `tivoli.tsm.devices.fcprte` to fail, and Fibre Channel support will not be enabled.

Once this is completed, using the `smitty devices` command will show a new menu entry, “Tivoli Storage Manager Devices”. Select **Tivoli Storage Manager Devices -> Fibre Channel SAN Attached Devices -> Discover Devices Supported by TSM**. If you have multiple Fibre Channel interfaces, you will be presented with a menu similar to the following:

```
 lqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqk
  x                             Parent FC Adapter                             x
  x                                                                           x
  x Move cursor to desired item and press Enter.                              x
  x                                                                           x
  x   fscsi0 Available 10-68-01 FC SCSI I/O Controller Protocol Device        x
  x   fscsi1 Available 20-58-01 FC SCSI I/O Controller Protocol Device        x
  x                                                                           x
  x F1=Help                 F2=Refresh              F3=Cancel                 x
  x Esc+8=Image             Esc+0=Exit              Enter=Do                  x
F1x /=Find                  n=Find Next                                       x
Esmqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqj
```

Select the correct adapter, and press `enter`. Another menu will be presented, asking you to select the parent FC adapter:

```
 Discover Devices Supported by TSM

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

                                              [Entry Fields]
  Parent FC Adapter                                  fscsi0
```

Select the correct entry, and press `enter`. After discovery completes, you will get a screen displaying the devices that have been found.

```
 COMMAND STATUS

Command: OK            stdout: yes           stderr: no

Before command completion, additional instructions may appear below.

====================== TSM Devices ====================
mt0 Available 10-68-01 Tape Drive
lb0 Available 10-68-01 Library/MediumChanger
mt1 Available 10-68-01 Tape Drive
```

At this point, the necessary drivers have been added to AIX, and the library is ready to be defined to Tivoli Storage Manager.

Windows 2000 preparation

Windows 2000 Build 2195 or greater is required. To determine the build of Windows 2000 that you have, right-click on **My Computer**, select **Manage**. The Computer Management window will appear; click on the plus sign (**+**) next to System Information, and click on **System Summary**. You should see something similar to Figure 24. Our lab system displayed version 5.0.2195 Build 2195.

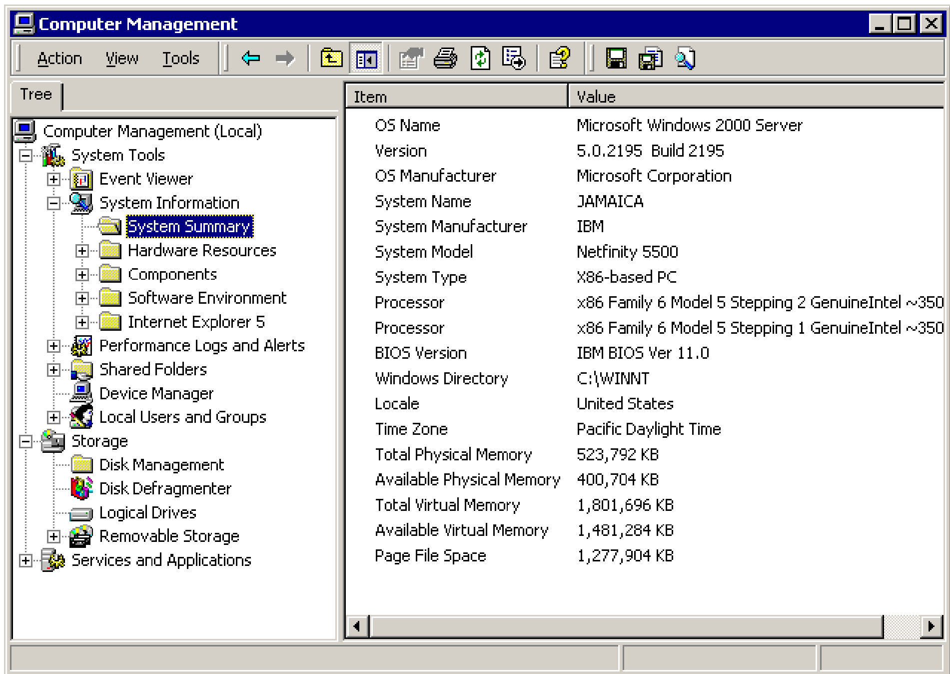

Figure 24. Determining Windows 2000 version

To implement library sharing on a Windows 2000 Server platform, it will probably be necessary to add drivers for the particular library that you are using. As we were using an IBM 3570 library for our test scenarios, this describes the implementation of that particular model.

After installing Windows 2000, ensure that the Fibre Channel interfaces are seen, and that the drivers are loaded. Right-click on **My Computer**, and select **Properties**. Then select the **Hardware** tab, and click on **Device Manager**.

There will be some devices that do not have drivers loaded, and these devices will be shown under the question mark (**?**) icon. Click on the plus sign (**+**) next to the question mark.

Click on the plus sign (**+**) next to SCSI and RAID controllers. You will see something similar to Figure 25.

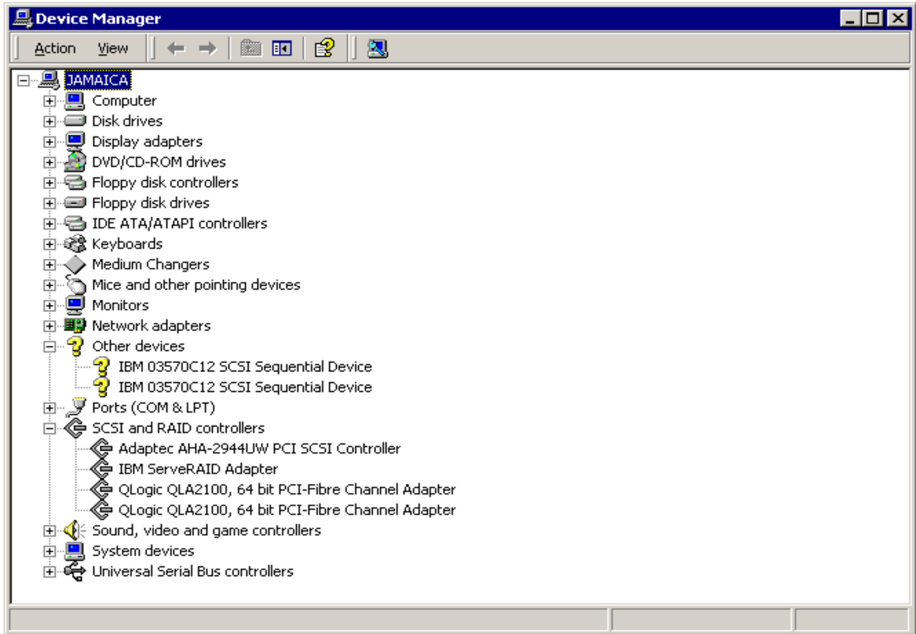

Figure 25. Windows 2000 SCSI controllers

This verifies that Windows 2000 has drivers loaded for the Qlogic Fibre Channel adapters, and that the library devices are seen.

For Windows 2000, you should **not** load the Windows `Atape` driver. Instead, use the `AdsmScsi` driver that is provided with Tivoli Storage Manager.

To use the driver, you must first load Tivoli Storage Manager. Install from your installation CD, using the normal Windows installation routine.

We recommend that you upgrade to the latest version of Tivoli Storage Manager, which is 4.1.1 as of this writing. The update for Tivoli Storage Manager 4.1.1 for Windows 2000 is the same as for Windows NT. You can obtain the update from:

```
ftp://service.boulder.ibm.com/storage/tivoli-storage-management/
maintenance/server/v4r1/NT/LATEST/
```

Note that the license files for the server are **not** included with the 4.1.1 update code. You must install the base 4.1.0 product in order to obtain the license files.

After installation of the Tivoli Storage Manager server code and any updates, launch the Tivoli Storage Manager Server Utilities. You will be prompted to begin the initial configuration wizard; you can cancel out of this wizard. You will see the screen shown in Figure 26.

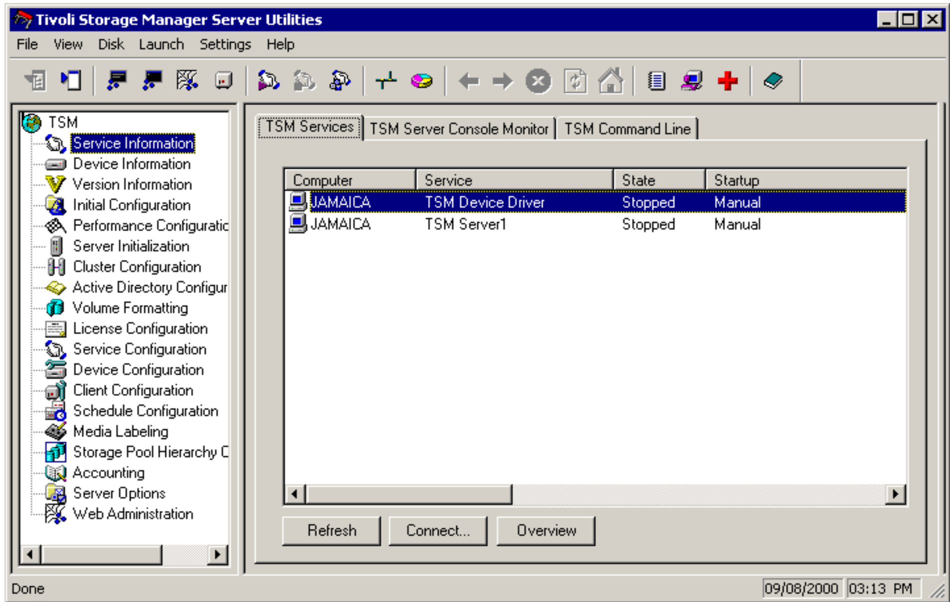

Figure 26. Tivoli Storage Manager Server Utilities

Right-click on the server instance shown in the right panel. Select **start** from the drop down menu; this will start the Tivoli Storage Manager server1 instance.

Note

Do **not** start the Tivoli Storage Manager Device Driver at this point!
The Device Configuration Wizard will attempt to start the device driver, and if it is started now, the Device Configuration Wizard will fail later.

Select Device Configuration in the left panel. You will be presented with the screen shown in Figure 27.

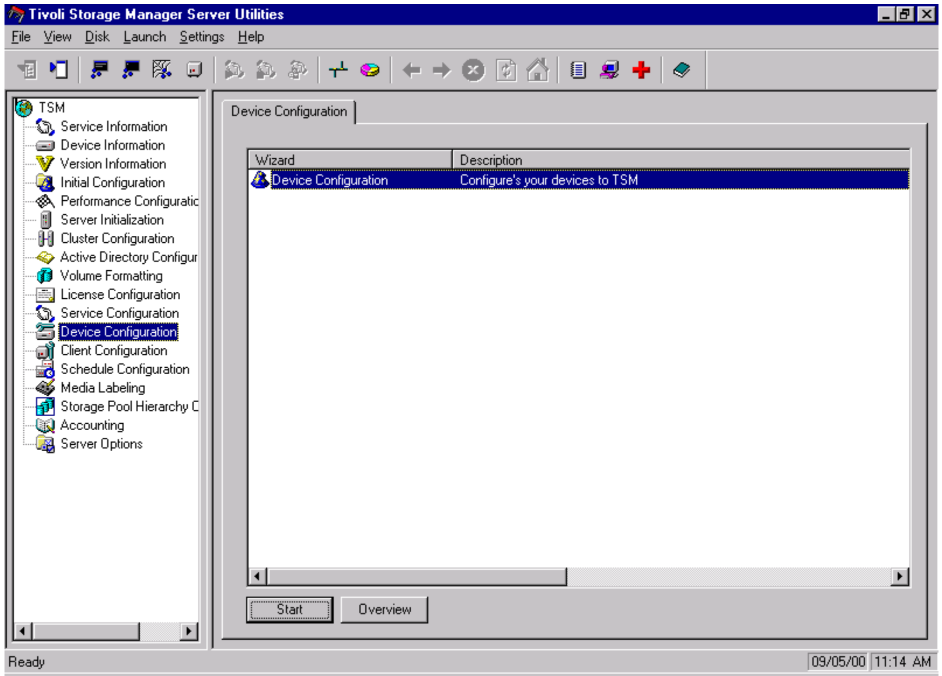

Figure 27. Start the Device Configuration Wizard

Click the **Start** button in the lower left corner of the right panel. This will start the Device Configuration Wizard.

Follow the prompts as the wizard presents them to you. Soon, you will be prompted to start the Tivoli Storage Manager Device Driver, as shown in Figure 28.

Figure 28. Starting the TSM device driver

After starting the driver, you are prompted to define which devices are available to Tivoli Storage Manager. Figure 29 shows the device selection window before selecting devices.

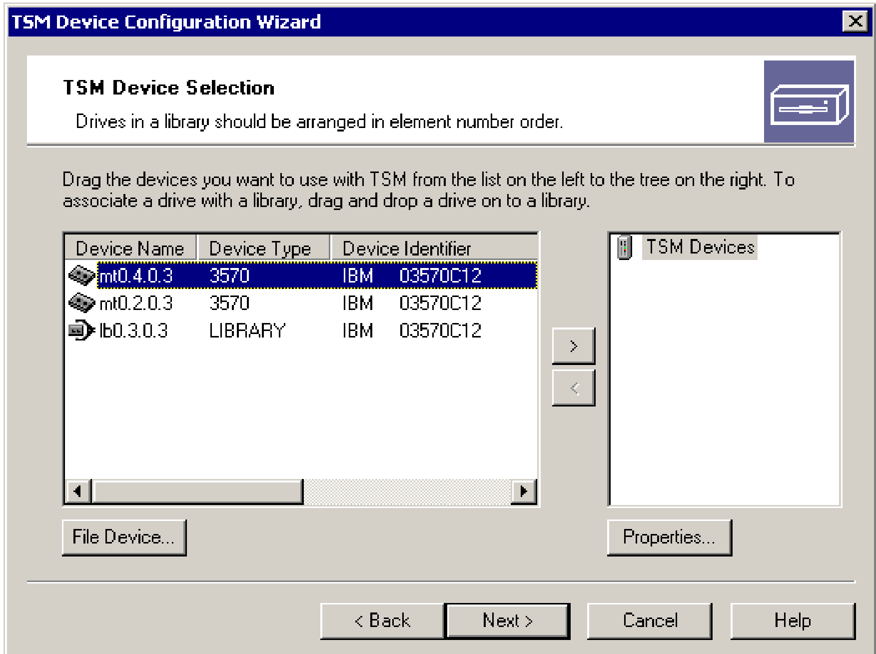

Figure 29. Device Selection — before assignment

Assign devices to Tivoli Storage Manager by dragging from the left panel, and dropping under the Tivoli Storage Manager Devices icon in the right panel.

Figure 30 shows the devices assigned to Tivoli Storage Manager. Note the alignment of devices under the library device.

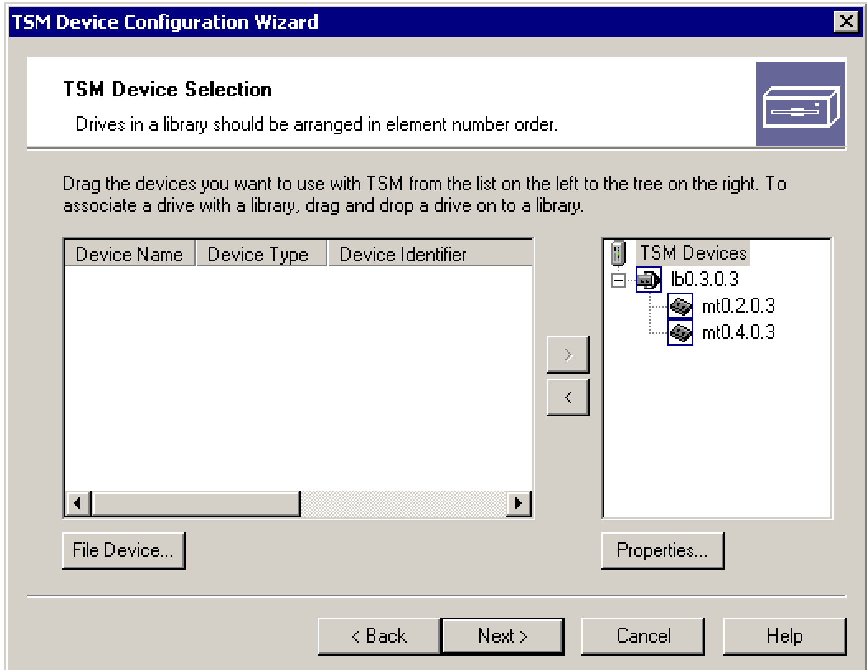

Figure 30. Device Selection — after assignment

Click the **Next** button, the wizard will complete, and you will be returned to the main window. Check your configuration by selecting **Device Information** in the left pane. The right pane will appear similar to Figure 31, showing your configuration.

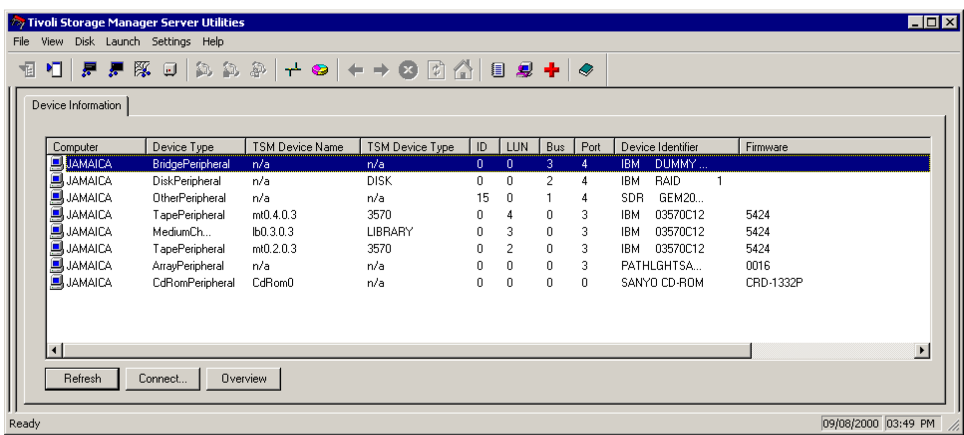

Figure 31. Device Configuration completed

With a Windows 2000 server, it is important to note that 3570 devices attached with this method, and having the Tivoli Storage Manager Device Driver loaded will still appear in the Windows Device Manager as not having drivers loaded. This is normal and expected.

Windows NT 4 preparation

In order to implement library sharing under Windows NT 4, you must be at Service Pack 5 or higher.

We recommend that you immediately upgrade the Tivoli Storage Manager server code to 4.1.1 upon installation. The base 4.1.0 code contains some installation glitches which are fixed by the 4.1.1 update. The 4.1.1 update code can be obtained from:

`ftp://service.boulder.ibm.com/storage/tivoli-storage-management/maintenanc`
`e/server/v4r1/NT/LATEST/`

Note that the server license files are **not** included with the 4.1.1 update package; you **must** install the base 4.1.0 code, then apply the 4.1.1 upgrade for the install to complete successfully.

Unlike an AIX installation, you should *not* load the Windows NT `Atape` driver, but rather use the `AdsmScsi` driver that comes with Tivoli Storage Manager. When you install Tivoli Storage Manager, this driver will be deposited on the system, and it is only necessary to start it through the device configuration wizard. The next section explains how to start the `AdsmScsi` driver.

After attaching the system to the Fibre Channel gateway and rebooting, start the Windows Control Panel, and select the **Tape Devices** icon. You will see a screen similar to Figure 32.

Note:

When you open the Tape Devices screen, you will be prompted to install drivers. **Do not** do so at this time. Install drivers using the Device Configuration wizard in Tivoli Storage Manager Server Utilities. See the following pages for further information on driver installation.

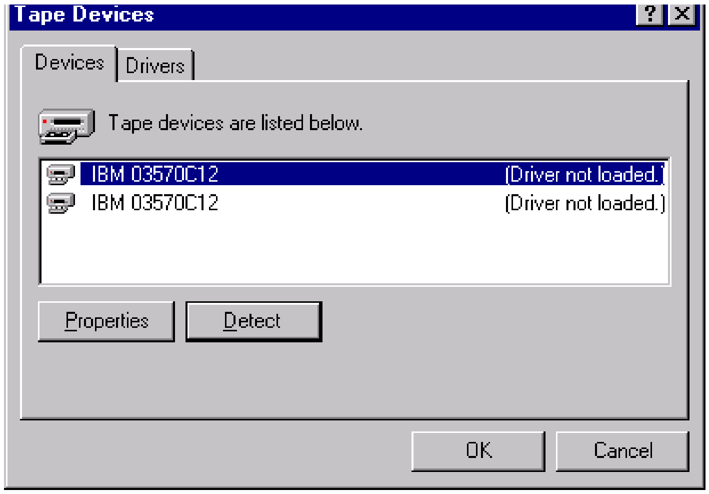

Figure 32. Windows NT Tape Devices window

After verifying that the tape devices can be “seen” by the operating system, close the Tape Devices window.

At this point, you can install the Tivoli Storage Manager code. Install the base code (Version 4.1), then immediately apply IP22102 to bring the server up to Version 4.1.1.

Launch the Tivoli Storage Manager Server Utilities, and select **Device Configuration** from the left pane as shown in Figure 33.

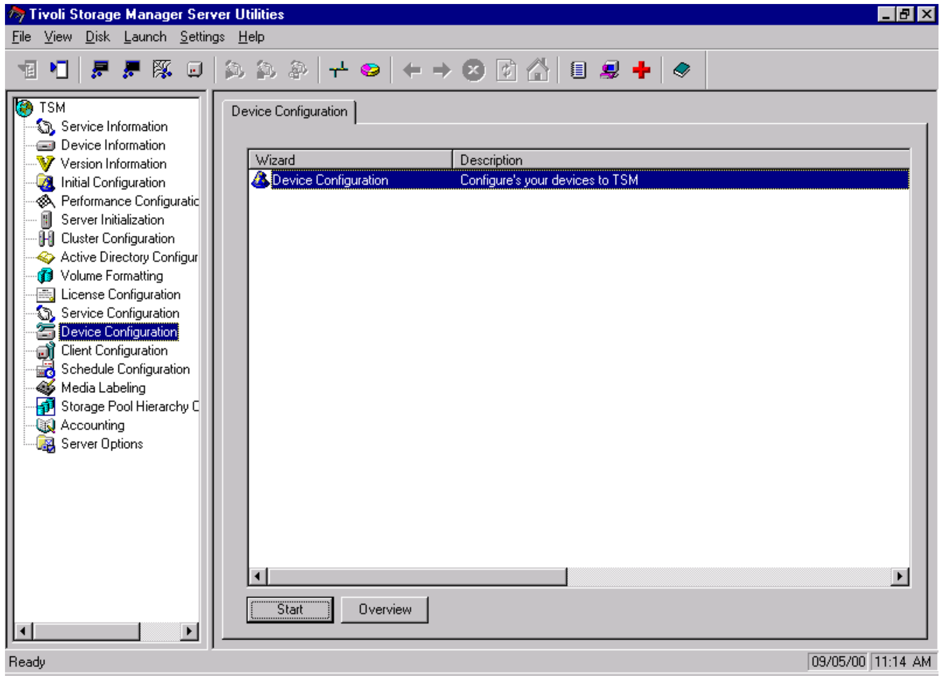

Figure 33. TSM Server Utilities Device Configuration Panel

Click on the **Start** button in the lower-left corner of the right pane. Follow the prompts presented. Soon, you will see a prompt to start the Tivoli Storage Manager Device Driver, as shown in Figure 34. Select **Yes** to start the Tivoli Storage Manager driver.

Figure 34. Starting the AdsmScsi device driver

Next, you will be prompted to define the devices to be used with Tivoli Storage Manager. Select the devices shown on the left pane, and drag them to the right pane, as shown in Figure 35. Note the positioning of the drive devices beneath the library device.

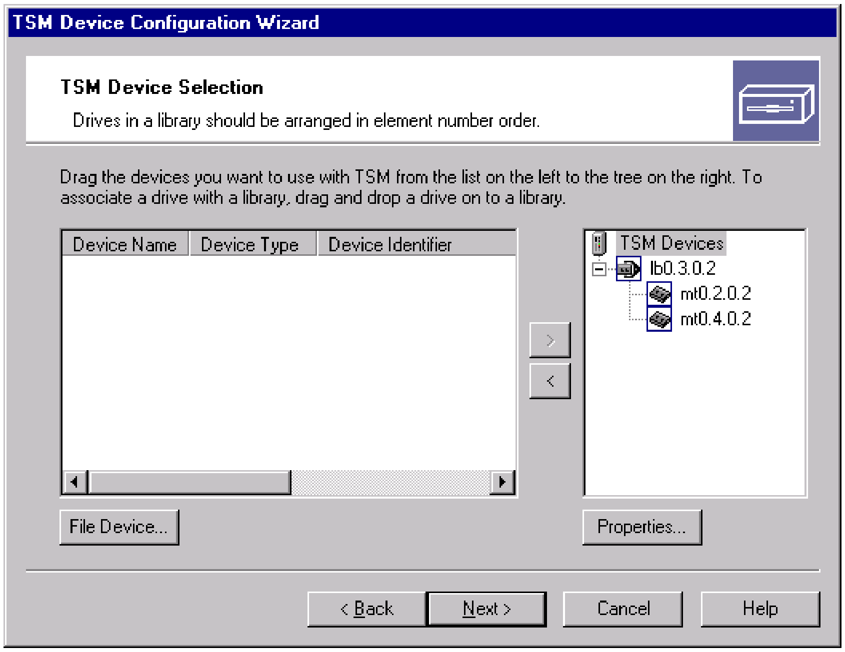

Figure 35. TSM Device Selection

After you have completed your device configuration, you can verify your selections by selecting **Device Information** from the main configuration screen. The device configuration should appear similar to Figure 36.

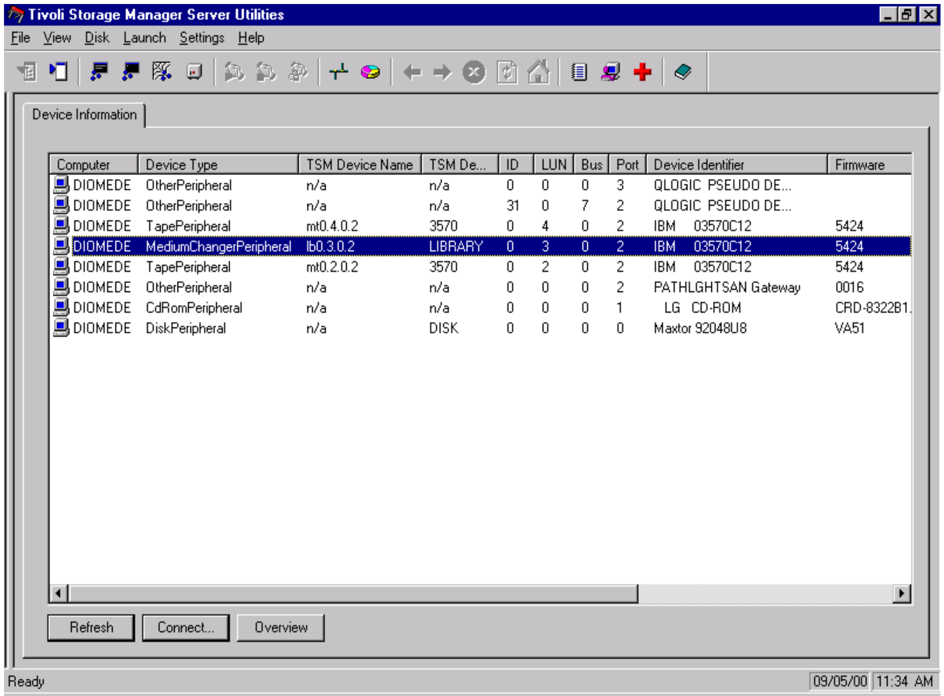

Figure 36. TSM Server Utilities device information window

At this point, the Tivoli Storage Manager configuration can proceed.

4.1.3 Using IBM SAN Data Gateway Router and Hub

The environment used for this section is shown in Figure 37.

For this first scenario, we shared a single SCSI tape library IBM 3570 C12 between two different Tivoli Storage Manager servers, using an IBM SAN Data Gateway Router, model 2108-R03, and an IBM Fibre Channel Hub, model 2103-H07. The two servers were:

- A Netfinity 5500 running Windows 2000 Server (Jamaica)
- A IBM PC system running Windows NT4 Server (Diomede)

We selected the Windows 2000 Server (Jamaica) to be the library manager.

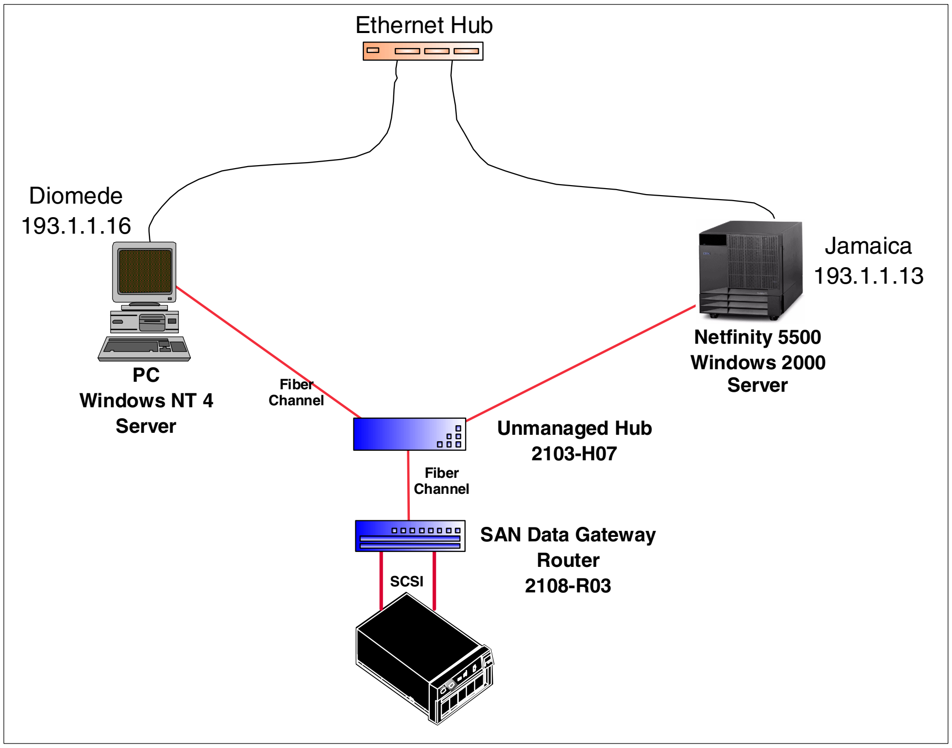

Figure 37. Using IBM SAN Data Gateway Router and Hub

Note that we have created an FC-AL (loop) configuration in this scenario by utilizing the 2103-H07 Hub. The two servers, Jamaica and Diomede, are in the Fibre Channel loop with the 2108-R03 Router; the total loop bandwidth is shared between these two servers.

4.1.3.1 IBM Fibre Channel Hub 2103-H07 preparation

Preparation of the Fibre Channel Hub is covered in Section 3.2.1, “Fibre Channel Hub 2107-H07 basic configuration” on page 54.

4.1.3.2 IBM SAN Data Gateway Router 2108-R03 preparation

Preparation of the SAN Data Gateway Router as we used it in this scenario is covered in Section 3.2.2, “SAN Data Gateway Router 2108-R03 basic configuration” on page 54.

After completing installing and preparing the hardware components, and completing the operating system preparation, you can proceed to Section

4.1.6, “Tivoli Storage Manager setup” on page 98 to configure Tivoli Storage Manager.

4.1.4 Using IBM SAN Data Gateway

The environment used for this section is shown in Figure 38.

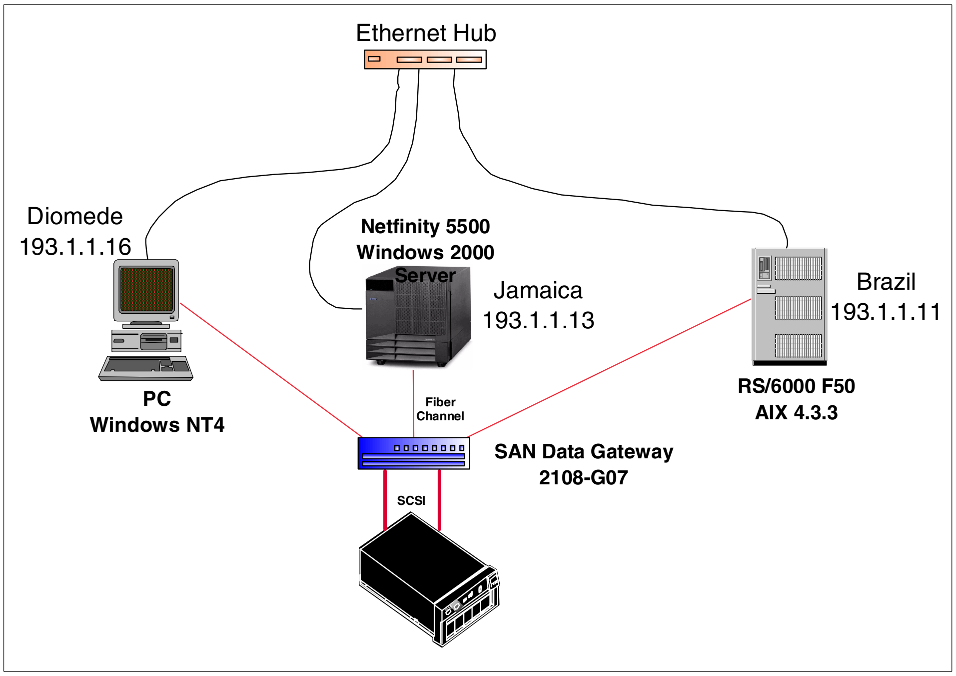

Figure 38. Using IBM SAN Data Gateway

For this scenario, we shared a single SCSI tape library between three different Tivoli Storage Manager servers, using an IBM SAN Data Gateway, model 2108-G07. The three servers were:

- An RS/6000 F50 running AIX 4.3.3 (Brazil)
- A Netfinity 5500 running Windows 2000 Server (Jamaica)
- A IBM PC running Windows NT4 Server (Diomede)

For this scenario we selected the RS/6000 F50 to be the library manager.

Note that in this scenario, we were running with our Fibre Channel ports set to point-to-point mode. The full bandwidth of the Fibre Channel connection from each server to the gateway is available for each node; connection bandwidth is not shared between the servers.

Preparation of the SAN Data Gateway as we used it for this scenario is covered in Section 3.2.3, “SAN Data Gateway 2108-G07 basic configuration” on page 62.

After completing installing and preparing the hardware components, and completing the operating system preparation, you can proceed to Section 4.1.6, “Tivoli Storage Manager setup” on page 98 to configure Tivoli Storage Manager.

4.1.5 Using IBM SAN Fibre Channel Switch and Data Gateway

The environment used for this section is shown in Figure 39.

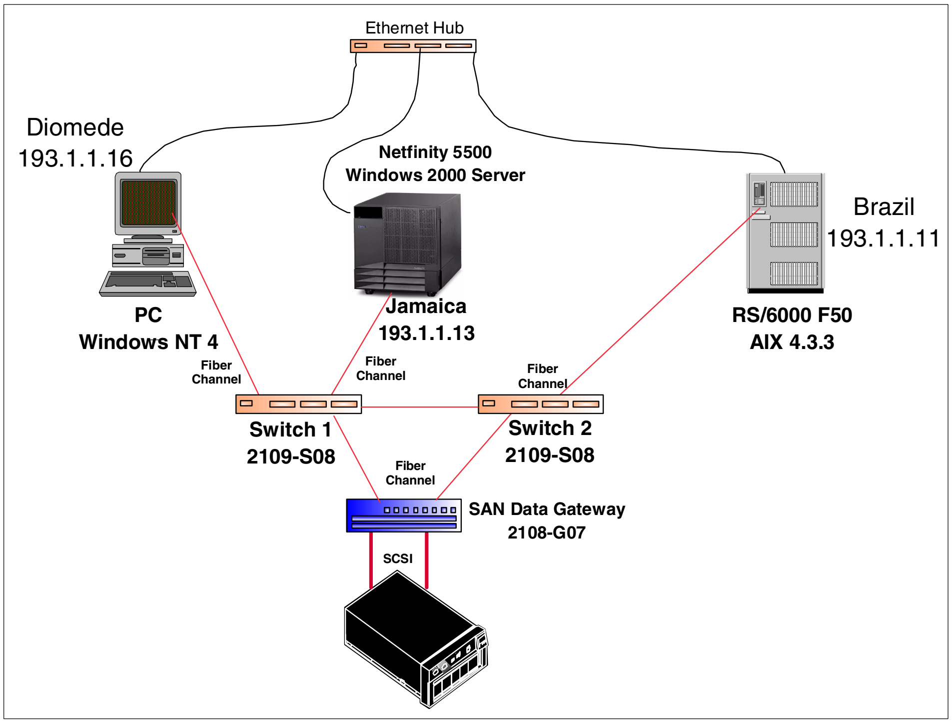

Figure 39. Using IBM SAN Fibre Channel Switch

Similar to the previous scenario we also shared a single SCSI tape library between three different Tivoli Storage Manager servers. We use now two IBM 2109-S08 switches to create our fabric, and an IBM SAN Data Gateway, model 2108-G07. The three servers were the same as in the previous scenario:

- An RS/6000 F50 running AIX 4.3.3 (Brazil)

- A Netfinity 5500 running Windows 2000 Server (Jamaica)
- A IBM PC running Windows NT4 Server (Diomede)

As before, we selected the RS/6000 F50 to be the library manager.

This scenario represents a switched fabric configuration.

Preparation of the SAN Fibre Channel Switch as we used it in this scenario is covered in Section 3.2, “Fibre Channel interconnect devices” on page 54. Configuration of the SAN Data Gateway device is the same as in the previous section.

4.1.6 Tivoli Storage Manager setup

In order to implement SCSI library sharing in a Tivoli Storage Manager environment, you will need to be at Release 3.7 or higher. As of Release 3.7, the only supported platforms are Windows NT and Sun Solaris, and only with a select number of tape libraries.

With the release of Tivoli Storage Manager 3.7.3, the supported servers now include AIX and Windows 2000. Support for additional SCSI libraries have been added as well. Information about currently supported servers and libraries is available at the following web page:

`http://www.tivoli.com/support/storage_mgr/san/libsharing.html`

For the purposes of this publication, we used Tivoli Storage Manager version 4.1.1.

The instructions in this section make the following assumptions:

- You have are able to “see” your library devices as described in the previous sections.
- You have loaded the drivers appropriate for your platforms, and the operating systems are configured to use those drivers.
- You have updated your hardware to a current microcode/firmware level.
- Tivoli Storage Manager is loaded on the servers, and has been updated to at least version 4.1.0.

Once the hardware, firmware, and operating systems are prepared as described in the previous sections, Tivoli Storage Manager configuration is virtually identical, regardless of the platform used. Accordingly, we will describe the process once for all three scenarios, as once the SAN configuration is in place, the Tivoli Storage Manager configuration and operation is the same, regardless of the SAN topology or platform used.

4.1.6.1 Configuring server-to-server communication

We began by configuring the library manager.

SCSI library sharing requires the use of server-server communication, so our first step is to enable that. Define the server name for the primary library manager, and set its password using the `SET SERVERNAME` and `SET SERVERPASSWORD` command:

```
tsm: BRAZIL>set servername brazil
ANR2094I Server name set to BRAZIL.

tsm: BRAZIL>set serverpassword brazil
ANR2131I Server password set.
```

In order to facilitate setting up server-server communication, allow crossdefine using the `SET CROSSDEFINE` command:

```
tsm: BRAZIL>set crossdefine on
ANR2135I Crossdefine set to ON.
```

Make sure that your library manager's high-level address, or IP address and low-level address, or port which is by default normally set to 1500 are set.

```
tsm: BRAZIL>SET SERVERHLADDRESS 193.1.1.11
ANR2132I Server hladdress set to 193.1.1.11.

tsm: BRAZIL>SET SERVERLLADDRESS 1500
ANR2133I Server lladdress set to 1500.
```

Repeat these steps on the Tivoli Storage Manager server which will be the library client, using the values appropriate for that server. In the lab, the commands on the library client were:

```
SET SERVER NAME JAMAICA
SET SERVERPASSWORD JAMAICA
SET CROSSDEFINE ON
SET SERVERHLADDRESS 193.1.1.13
SET SERVERLLADDRESS 1500
```

At this point, you can define each server to the other using the `DEFINE SERVER` command. As we have enabled crossdefine, this will create the server definitions on both the library manager (shown below):

```
tsm: BRAZIL>define server jamaica serverpa=jamaica hladdress=193.1.1.13 lladdress=1500
crossdefine=yes
ANR1660I Server JAMAICA defined successfully.
tsm: BRAZIL>q server

Server   Comm.  High-level    Low-level   Days Server   Virtual  Allow
Name     Method Address       Address    Since Password Volume   Replacement
                                          Last Set      Password
                                        Access          Set
-------- ------ ------------- --------- ------ -------- -------- -----------
DIOMEDE  TCPIP  193.1.1.16    1500          14 Yes      Yes      No
JAMAICA  TCPIP  193.1.1.13    1500          <1 Yes      No       No
```

And on the library client:

```
ANS8000I Server command: 'q server'

                Server Name: BRAZIL
               Comm. Method: TCPIP
         High-level Address: 193.1.1.11
          Low-level Address: 1500
     Days Since Last Access: <1
        Server Password Set: Yes
Virtual Volume Password Set: Yes
          Allow Replacement: No


ANS8002I Highest return code was 0.
```

Now that the servers are known to each other, define the library on each server.

4.1.6.2 Defining the shared SCSI library

This section describes how to define the shared library to the Tivoli Storage Manager servers acting as library manager or library client.

Important

When defining a shared SCSI library, you define the library as “shared” on the library manager, but drives are defined normally. There is no “shared” parameter that applies to individual devices within a library.

When defining the library on a library client,use the special parameters PRIMARYLIBMANAGER and LIBTYPE=SHARED to tell the client which server is the library manager. Note that you still use the same drive devices that you would normally use.

Both the library name and the drive names *must* match on the library manager and library clients!

Begin with the library manager:

```
tsm: BRAZIL>define library 3570lib libtype=scsi device=/dev/rmt1.smc shared=yes
ANR8400I Library 3570LIB defined.

tsm: BRAZIL>define drive 3570lib drive0 device=/dev/rmt1 element=16
ANR8404I Drive DRIVE0 defined in library 3570LIB.

tsm: BRAZIL>define drive 3570lib drive1 device=/dev/rmt3 element=17
ANR8404I Drive DRIVE1 defined in library 3570LIB.
```

After defining the library and drive devices, issue a `QUERY LIBRARY F=D` command to check the definition of the library:

```
tsm: BRAZIL>q libr f=d

                  Library Name: 3570LIB
                  Library Type: SCSI
                        Device: /dev/rmt1.smc
              Private Category:
              Scratch Category:
              External Manager:
                        Shared: Yes
       Primary Library Manager:
     Secondary Library Manager:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 09/19/00   14:47:03
```

Note that the library appears as “shared” in the listing above.

Now that the library is defined and shared on the library manager, define the library and drives to the library client Tivoli Storage Manager servers. Define the library using the PRIMARYLIBMANAGER and library type of SHARED parameters:

```
tsm: JAMAICA>define libr 3570lib libtype=shared primarylibmanager=brazil
ANR8400I Library 3570LIB defined.
```

Query the library from the client. Notice that the library appears as **not** shared from the library client’s perspective, but does indicate that there is a library manager (brazil).

```
tsm: JAMAICA>q libr f=d

                  Library Name: 3570LIB
                  Library Type: SHARED
                        Device:
              Private Category:
              Scratch Category:
              External Manager:
                RSM Media Type:
                        Shared: No
       Primary Library Manager: BRAZIL
     Secondary Library Manager:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 09/20/2000 15:54:29
```

Now, define the drives, remembering to use the same names for the drives and library as on the library manager server. Notice that it is not necessary to specify the element number of the drive, as the library manager handles all requests for tape movement:

```
tsm: JAMAICA>define drive 3570lib drive0 device=mt0.4.0.3
ANR8404I Drive DRIVE0 defined in library 3570LIB.

tsm: JAMAICA>define drive 3570lib drive1 device=mt0.2.0.3
ANR8404I Drive DRIVE1 defined in library 3570LIB.
```

Define storage pools and device class definitions as you would normally.

This completes the shared SCSI library setup for Tivoli Storage Manager.

4.1.7 Using a shared SCSI library with Tivoli Storage Manager

When using a shared SCSI library, it is important to remember that the library clients cannot directly access the library robotics themselves. Any requests for tape movement must be requested through the library manager.

While library clients do not directly access the library robotics, once a volume is mounted in a drive, the library clients are able to write or read data directly to the drive through the SAN data path.

4.1.7.1 Volume query commands

On the library client, querying the contents of the library will not report the library contents:

```
tsm: JAMAICA>q libvolume
ANR2034E QUERY LIBVOLUME: No match found using this criteria.
ANS8001I Return code 11.
```

While on the server, the same command will display the contents of the library, as well as which Tivoli Storage Manager server ”owns” the volumes:

```
tsm: BRAZIL>q libvolume
Session established with server BRAZIL: AIX-RS/6000
  Server Version 4, Release 1, Level 1.0
  Server date/time: 09/21/00   13:15:16  Last access: 09/21/00   11:56:41


Library Name   Volume Name   Status       Owner        Last Use    Home Element
------------   -----------   ----------   ----------   ---------   ------------
3570LIB        080CED        Private      BRAZIL                   40
3570LIB        083332        Scratch                               41
3570LIB        083964        Private      BRAZIL       Data        42
3570LIB        083983        Private      BRAZIL       DbBackup    35
3570LIB        085135        Scratch                               37
3570LIB        08527E        Scratch                               33
3570LIB        085E0A        Scratch                               34
3570LIB        085EA7        Private      JAMAICA      Data        43
3570LIB        085EA8        Private      JAMAICA      DbBackup    36
3570LIB        086094        Private      BRAZIL                   38
3570LIB        086247        Private      JAMAICA      Data        39
```

Querying the volumes on any attached server will report only those volumes that belong to that server.

For the library client:

```
tsm: BRAZIL>q vol

Volume Name               Storage      Device      Estimated    Pct   Volume
                          Pool Name    Class Name   Capacity   Util   Status
                                                        (MB)
------------------------  -----------  ----------  ---------  -----  --------
083964                    TAPEPOOL     TEST          5,000.0   54.2  Filling

tsm: JAMAICA>q volhist type=dbbackup

       Date/Time: 09/21/2000 11:39:55
     Volume Type: BACKUPFULL
   Backup Series: 1
Backup Operation: 0
      Volume Seq: 1
    Device Class: 357CLASS1
     Volume Name: 085EA8
 Volume Location:
         Command:
```

For the server:

```
tsm: BRAZIL>q libvolume
Session established with server BRAZIL: AIX-RS/6000
  Server Version 4, Release 1, Level 1.0
  Server date/time: 09/21/00   13:15:16  Last access: 09/21/00   11:56:41

Library Name   Volume Name   Status       Owner        Last Use    Home Element
------------   -----------   ----------   ----------   ---------   ------------
3570LIB        080CED        Private      BRAZIL                   40
3570LIB        083332        Scratch                               41
3570LIB        083964        Private      BRAZIL       Data        42
3570LIB        083983        Private      BRAZIL       DbBackup    35
3570LIB        085135        Scratch                               37
3570LIB        08527E        Scratch                               33
3570LIB        085E0A        Scratch                               34
3570LIB        085EA7        Private      JAMAICA      Data        43
3570LIB        085EA8        Private      JAMAICA      DbBackup    36
3570LIB        086094        Private      BRAZIL                   38
3570LIB        086247        Private      JAMAICA      Data        39

tsm: BRAZIL>q volhist type=dbb

       Date/Time: 09/21/00   10:33:36
     Volume Type: BACKUPFULL
   Backup Series: 5
Backup Operation: 0
      Volume Seq: 1
    Device Class: 3570CLASS
     Volume Name: 083983
 Volume Location:
         Command:
```

4.1.7.2 Auditing the library

Performing an AUDIT LIBRARY command from the library client does not actually perform a physical audit on the shared library; instead, the library client's database is checked against the library manager's database:

```
tsm: JAMAICA>audit library 3570lib
ANS8003I Process number 4 started.
```

The activity log on the library client will show entries similar to the following. Note that during this audit, no tape mounts were performed:

```
ANR0984I Process 4 for AUDIT LIBRARY started in the BACKGROUND at 11:44:49.
ANR8457I AUDIT LIBRARY: Operation for library 3570LIB started as process 4.
ANR8461I AUDIT LIBRARY process for library 3570LIB completed successfully.
ANR0985I Process 4 for AUDIT LIBRARY running in the BACKGROUND completed with
completion state SUCCESS at 11:44:49.
```

Performing a library audit on the library manager mounts or inspects tapes in the library, as it does in non-shared configurations:

```
tsm: BRAZIL>audit library 3570lib
ANS8003I Process number 8 started.
```

Activity log output on the server appears as below:

```
09/21/00   13:36:08  ANR2017I Administrator ADMIN issued command: AUDIT LIBRARY
                      3570lib
09/21/00   13:36:08  ANR0984I Process 8 for AUDIT LIBRARY started in the
                      BACKGROUND at 13:36:08.
09/21/00   13:36:08  ANR8457I AUDIT LIBRARY: Operation for library 3570LIB
                      started as process 8.
09/21/00   13:41:37  ANR8461I AUDIT LIBRARY process for library 3570LIB
                      completed successfully.
09/21/00   13:41:37  ANR0985I Process 8 for AUDIT LIBRARY running in the
                      BACKGROUND completed with completion state SUCCESS at
                      13:41:37.
```

Notice that several minutes have elapsed as the audit was performed; during that time, the library was mounting and reading tapes.

4.1.7.3 Checkin and checkout of tape volumes

Checking in or out of a volume from a shared library must be done from the library manager.

Attempting to check a tape out of a shared library from a library client receives the following error message:

```
tsm: JAMAICA>checkout libvolume 3570lib 086247
ANR8403E CHECKOUT LIBVOLUME: Operation not allowed for SHARED libraries.
ANS8001I Return code 3.
```

However, checking a tape out from the library manager succeeds:

```
tsm: BRAZIL>checkout libvolume 3570lib 085ea7
ANS8003I Process number 9 started.

tsm: BRAZIL>q req
ANR8352I Requests outstanding:
ANR8322I 001: Remove 3570 volume 085EA7 from entry/exit port of library
3570LIB; issue 'REPLY' along with the request ID when ready.

tsm: BRAZIL>reply 001
ANR8499I Command accepted.
```

If a tape is checked out of a shared SCSI library, and a node which is a client of a library client requires that tape, a single message is seen on the library client, and mount request prompts appear on the library manager.

In our scenario, the backup/archive client on Jamaica requires a volume for a restore process. Jamaica initiates a session with its server (which in this case is also running on Jamaica), and requests the data. The Jamaica server sends a request to its library manager (Brazil) requesting the tape mount.

The activity log on the library client (Jamaica) shows the following entries:

```
ANR0400I Session 158 started for node JAMAICA (WinNT) (Named Pipe).
ANR4709W Intervention required to mount volume 085EA7 on Library Manager Server
BRAZIL.
```

The activity log on the server shows the mount request, and operator responses when the tape is mounted:

```
ANR0408I Session 285 started for server JAMAICA (Windows NT) (Tcp/Ip) for
library sharing.
ANR8308I 002: 3570 volume 085EA7 is required for use in library 3570LIB;
CHECKIN LIBVOLUME required within 10 minutes.

ANR2017I Administrator ADMIN issued command: CHECKIN libvolume 3570lib 085ea7
status=scratch
ANR0984I Process 10 for CHECKIN LIBVOLUME started in the BACKGROUND at
14:00:42.
ANR8422I CHECKIN LIBVOLUME: Operation for library 3570LIB started as process
10.
ANR8323I 003: Insert 3570 volume 085EA7 R/W into entry/exit port of library
3570LIB within 60 minute(s); issue 'REPLY' along with the request ID when
ready.
ANR8308I 002: 3570 volume 085EA7 is required for use in library 3570LIB;
CHECKIN LIBVOLUME required within 8 minutes.
ANR2017I Administrator ADMIN issued command: QUERY REQ
ANR8352I Requests outstanding:
ANR8323I 003: Insert 3570 volume 085EA7 R/W into entry/exit port of library
3570LIB within 60 minute(s); issue 'REPLY' along with the request ID when
ready.
ANR2017I Administrator ADMIN issued command: REPLY 3
ANR8499I Command accepted.
ANR8335I 003: Verifying label of 3570 volume 085EA7 in drive DRIVE0
(/dev/rmt1).

ANR8328I 003: 3570 volume 085EA7 mounted in drive DRIVE0 (/dev/rmt1).
ANR8427I CHECKIN LIBVOLUME for volume 085EA7 in library 3570LIB completed
successfully.
ANR8427I CHECKIN LIBVOLUME for volume 085EA7 in library 3570LIB completed
successfully.
ANR0985I Process 10 for CHECKIN LIBVOLUME running in the BACKGROUND completed
with completion state SUCCESS at 14:01:29.
ANR8337I 3570 volume 085EA7 mounted in drive DRIVE1 (/dev/rmt3).
ANR0409I Session 285 ended for server JAMAICA (Windows NT).
```

4.1.7.4 MOVE MEDIA and MOVE DRMEDIA commands

The `MOVE MEDIA` and `MOVE DRMEDIA` commands are handled in the same manner as checkin/checkout requests. Requests for tape handling operations are sent to the library manager, which handles the requests.

4.1.7.5 QUERY MOUNT command

The `QUERY MOUNT` command behaves somewhat differently on the library client than it does on the library manager. On the library manager, it behaves as it always has, showing volumes mounted in drives, regardless of the shared library client which owns that particular volume. The screen below shows a volume which was mounted by a library client (Jamaica), which is now idle:

```
tsm: BRAZIL>q mount
ANR8329I 3570 volume 085EA7 is mounted R/W in drive DRIVE1 (/dev/rmt3), status:
IDLE.
ANR8334I         1 volumes found.
```

If the same command is run on the library client, which actually requested the mount in the first place, the following output is seen:

```
tsm: JAMAICA>q mount
Session established with server JAMAICA: Windows NT
  Server Version 4, Release 1, Level 1.0
  Server date/time: 09/21/2000 14:16:14  Last access: 09/21/2000 13:46:06

ANR2034E QUERY MOUNT: No match found using this criteria.
```

While the volume is in active use, or before the MOUNTRETENTION period on the library client expires, issuing a `q mount` command on the library client will show the mounted volume; however, in no case will it show volumes mounted by other library clients, or the library manager.

We discovered that the amount of time that the library client “sees” a volume as being mounted is governed by the MOUNTRETENTION parameter for the device class for that device on the library client. It does not, however, control physically dismounting the volume in the library. The physical dismount of the volume is governed by the MOUNTRETENTION parameter on the library manager.

4.1.7.6 Labeling volumes

In a shared SCSI library environment, tape labeling must be performed on the library manager. Attempting to perform a labeling operation from a library client gives the following error:

```
tsm: JAMAICA>label libvolume 3570lib 080ced checkin=scratch overwrite=yes
Session established with server JAMAICA: Windows NT
  Server Version 4, Release 1, Level 1.0
  Server date/time: 09/21/2000 15:12:17  Last access: 09/21/2000 14:16:14

ANR8494E LABEL LIBVOLUME: An option specified is not valid for SHARED
libraries.
ANS8001I Return code 3.
```

Aside from this requirement, media labeling operations are the same in a shared environment as in a non-shared environment.

4.1.7.7 Scratch media sharing with shared SCSI libraries

With shared SCSI libraries, it is now possible to share like-scratch media between library clients. This is due to the fact that the Tivoli Storage Manager server selected as library manager serves as a central repository for information about volumes contained in the library. Since all media handling requests are handled by the library manager, the library manager can keep track of which tape belongs to which library client, and which tapes are “free” and available for use by any library client.

There is no special configuration required to utilize this feature; by simply keeping an adequate supply of labelled scratch tapes in the library, the library manager can allocate them as needed to the requesting library clients.

Once a particular volume is used by a library client, the library manager flags that volume as belonging to that client, and prevents other library clients from accessing it.

When all the data on a volume has been expired, the volume is returned to the scratch state on the library manager.

4.1.7.8 Library resource contention

While the SCSI library robotics are controlled solely by the library manager, the drives within the library are seen by all library clients. So how is contention for drive access handled?

When two library clients attempt to access the same drive(s) in a shared library, the library client that reserves the drive first will obtain control of the drive. The other library client will wait for access for the time limit specified by the MOUNTWAIT parameter.

This is not true if the second library client has a higher priority process. For example, if the first library client is accessing the drive for the purposes of a client backup, and the second library client requires the drive in order to do a database backup, the first library client will suspend its operation and relinquish the drive to the higher priority process.

4.2 Sharing an IBM 3494 library

This section covers the setup of a SAN-attached 3494 library which is to be shared between servers in a heterogeneous environment.

A detailed explanation of an IBM 3494 configuration and use is beyond the scope of this publication. Accordingly, we will confine our discussion to those elements of an IBM 3494 configuration that pertain to its implementation in a SAN attached environment.

4.2.1 Overview of IBM 3494 architecture and operation

The IBM 3494 tape library is a high capacity automated tape library, with the ability to handle up to 6,240 tape cartridges and 62 tape drives. A more comprehensive overview of the Magstar 3494 Tape Library can be found in the publication *Magstar 3494 Library Introduction and Planning Guide*, GA32-0279.

The IBM 3494 can be attached to and addressed by multiple hosts concurrently. Access to individual volumes is controlled by the use of volume categories, which are assigned to each individual cartridge by the library manager and host application. For Tivoli Storage Manager, the default categories assigned are 300 (for private volumes) and 301 (for scratch volumes). As we will discuss in Section 4.2.4, “Tivoli Storage Manager configuration” on page 118, additional Tivoli Storage Manager servers attached to the library will need to use different volume categories.

The IBM 3494 library differentiates between the logical data path and the logical control path. Communication to the library for these two paths can be implemented together on the same channel, as is possible when using ESCON or parallel connections to the library, or separately, as is the case when using Fibre Channel for the data path and an ethernet connection for control.

Figure 40 shows different attachment alternatives of a IBM 3494 tape library to hosts. It shows and attachment via a parallel channel. In this case, both data and control traffic flow via the parallel connection. The second alternative shows a host attached to a 3494 using separate data and control paths. In this instance, library control is managed through the an RS232 or an TCP/IP connection, while data flowing to a tape device flows through the SCSI attachment.

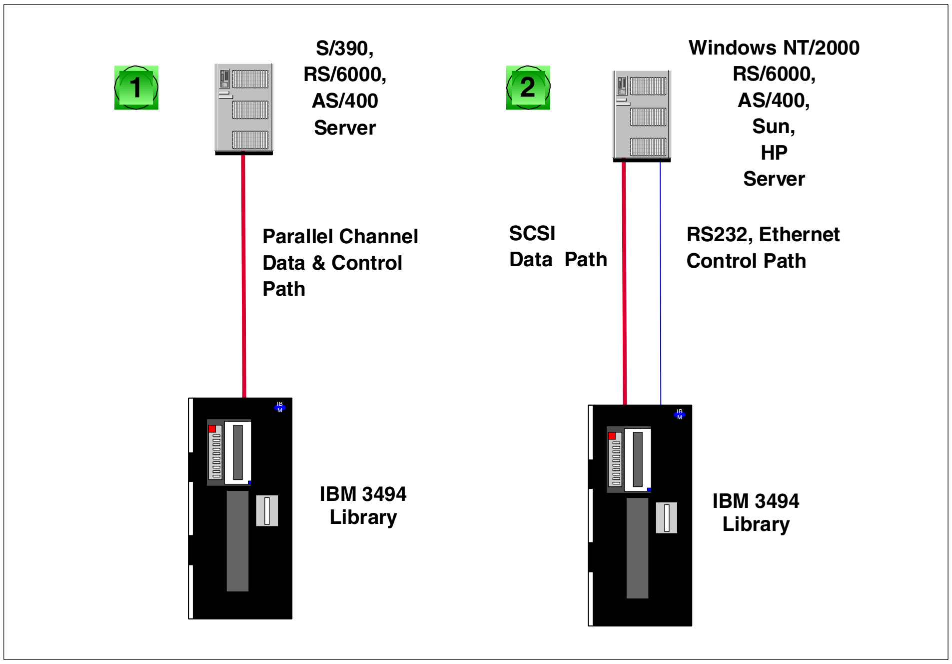

Figure 40. IBM 3494 tape library — connection alternatives

Figure 41 shows a representation of the communication paths internal to the 3494 for two servers using SCSI and parallel attachment methods.

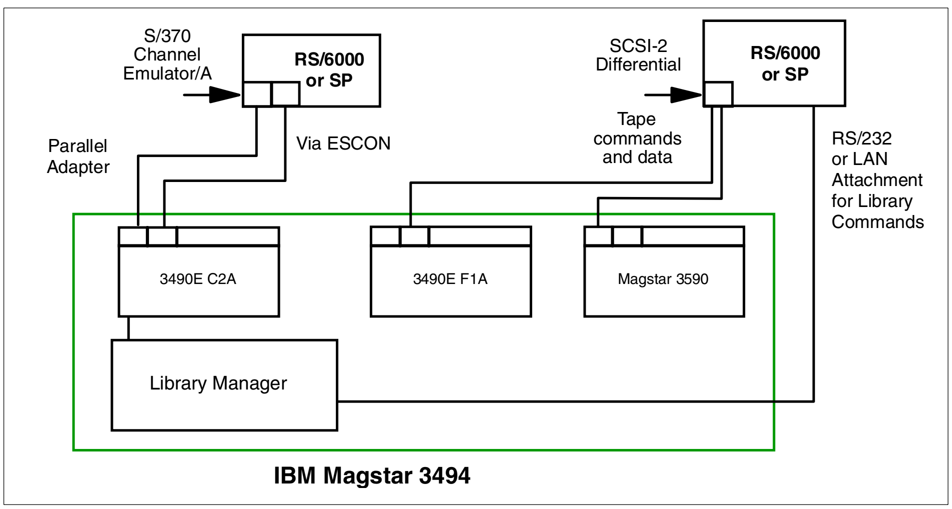

Figure 41. IBM 3494 library — communication paths

Refer to the following publications for further information on these topics:

- *Magstar 3494 Tape Library Introduction and Planning Guide*, GA32-0279
- *Magstar 3494 Tape Library Operator Guide*, GA32-0280
- *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632
- *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409

4.2.2 Terminology: partitioning versus sharing

Due to the enhanced capabilities of the 3494 library, the term “sharing” is used differently in the 3494 environment than in a SCSI library environment. This section attempts to provide an overview of these differences.

Refer to Chapter 4 of the publication *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409, for further details.

Note:

In this section, we differentiate between “host” and “application”, as it is possible to have different applications running on the same host, which use different volume categories, and therefore use a completely different set of tape volumes in the same library.

Sharing

Sharing in the 3494 world refers to multiple hosts having the ability to access a common set of volumes in a given library. An example of this would be two MVS systems accessing a common set of tapes in a library. In a 3494 shared environment, the attached hosts and applications use the same volume categories. The 3494 library manager does not manage whether or not a host has access to a volume; it is up to the host applications to determine how access is managed.

In a true shared 3494 environment, care must be taken to ensure that the accesses from different hosts will not conflict with each other. If the shared access is not adequately planned, data corruption could result.

Partitioning

Partitioning refers to dividing a physical library into multiple logical libraries. In a partitioned configuration, each attached host can only access the volumes in its partition. A volume which is assigned to a one host application cannot be (easily) accessed by a different host application.

Library volumes are assigned different volume categories in a partitioned 3494. Different hosts (or possibly different applications) “own” a particular volume category or categories. The 3494 library manager will only allow the application which references that category to access those volumes with that category.

It is still important to ensure that the host application does not attempt to use a category which is already allocated to another application.

Regardless of whether the library accessed is shared or partitioned, individual drives can be addressed by only one system at a time.

For the purposes of this publication, what we refer to as “3494 sharing” is actually what has been referred to as “3494 partitioning” in the past.

4.2.3 Setting up IBM 3494 tape library sharing in a SAN

In this section we are documenting the setup of an IBM 3494 tape library shared between an Tivoli Storage Manager Server running on AIX and a Tivoli Storage Manager server running on Windows 2000. The detailed layout is shown in Figure 42.

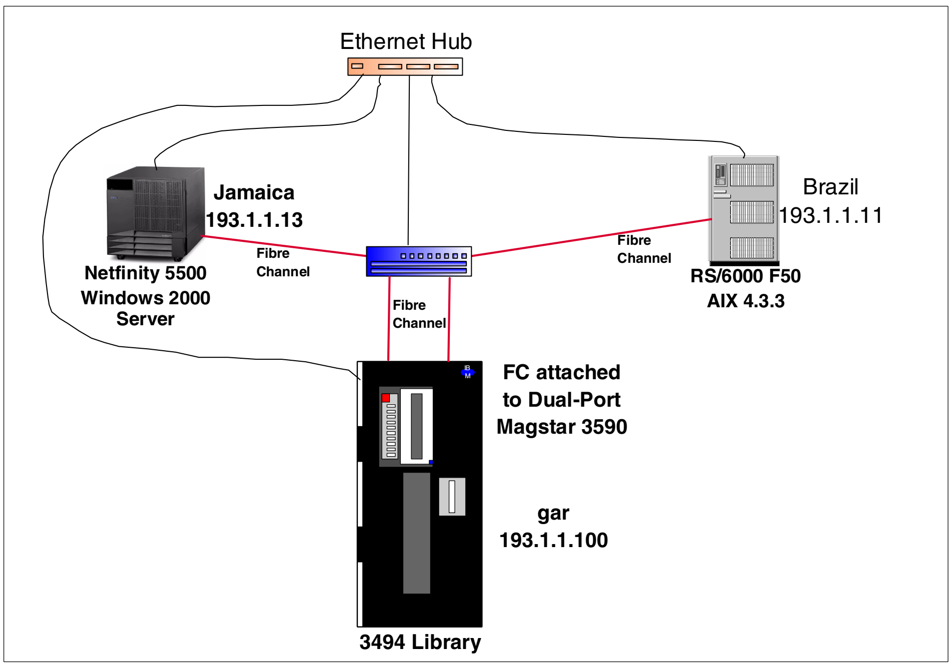

Figure 42. Lab setup for IBM 3494 partitioning

The IBM 3494 is equipped with native Fibre Channel attached IBM 3590 tape drives and the drives are direct connected to the server systems using an IBM SAN Fibre Channel switch 2109-S08. The library control path is utilizing an Ethernet based LAN.

The following Subsection will cover the following steps needed to setup IBM 3494 to be shared in a SAN environment.

1. Prepare 3494 library (Section 4.2.3.1, “3494 library preparation” on page 114)
2. Configure SAN environment (Section 4.2.3.2, “SAN component preparation” on page 114)
3. Prepare server hardware (Section 3.1, “Host Bus Adapters (HBAs)” on page 51)
4. Prepare server operating system and load drivers (section 4.2.3.4, “Operating system preparation” on page 114)
5. Perform Tivoli Storage Manager configuration. (section 4.2.4, “Tivoli Storage Manager configuration” on page 118)

4.2.3.1 3494 library preparation

Preparation of the 3494 library is identical to a dedicated configuration, with the exception of physical connections, and the requirement of adding the names and IP addresses of any additional connected hosts to the library manager’s database.

If the hosts which will be sharing the library are not added to the library manager’s database, the library manager will not allow access.

4.2.3.2 SAN component preparation

Refer to section 3.2.4, “IBM Fibre Channel Switch 2109-S08 basic configuration” on page 69 for information on switch setup as it was used in this scenario.

4.2.3.3 System hardware preparation

Preparation of the server host bus adapters (HBAs) is the same for the 3494 tape library sharing as for a SCSI tape library sharing using Tivoli Storage Manager. Refer to Section 3.1, “Host Bus Adapters (HBAs)” on page 51 for information regarding supported HBAs and firmware updates.

4.2.3.4 Operating system preparation

This section deals with preparation of the operating system environment for the lab scenario.

AIX preparation

Preparation of an RS/6000 AIX server to participate in a partitioned 3494 environment is almost the same as for a shared SCSI library, with the exception that the `atldd` driver (also known as the `lmcpd` driver) will also needed to be installed. The driver is required to communicate with the 3494 library manager. Refer to Section “AIX preparation” on page 80 for prerequisites, and installation of the `Atape` driver on AIX.

The `atldd` driver for AIX can be obtained from:

`ftp://index.storsys.ibm.com/devdrvr/AIX/`

The latest version, as of this writing, is 4.1.5.0.

Download the `atldd` driver to an empty directory on the server to be attached to the 3494. Run `inutoc` command to create the `.toc` file in that directory:

```
root@brazil:/tmp/atldd >ls -l
total 1600
-rw-r--r--   1 root     system    819200 Sep 15 13:39 atldd.4.1.5.0.bin
root@brazil:/tmp/atldd >
root@brazil:/tmp/atldd >inutoc .
root@brazil:/tmp/atldd >
root@brazil:/tmp/atldd >ls -l
total 1608
-rw-r--r--   1 root     system       141 Sep 15 13:43 .toc
-rw-r--r--   1 root     system    819200 Sep 15 13:39 atldd.4.1.5.0.bin
root@brazil:/tmp/atldd >
```

Then use `smitty install` command to install the `atldd` driver. When it completes, check to verify that the install completed sucessfully:

```
root@brazil:/tmp/atldd >lslpp -l all | grep atldd
  atldd.driver               4.1.5.0  COMMITTED  IBM Automated Tape Library
```

Once the driver is loaded, the configuration file `/etc/ibmatl.conf` must be edited to point to the library manager. The figure below shows the single line entry that was require to talk to our lab 3494:

```
root@brazil:/tmp/atldd >tail -1 /etc/ibmatl.conf
gar    193.1.1.100    gar
```

The entry above defines a library called “gar” at IP address 193.1.1.100, which will be known as “gar” to the host system.

Verify that the server can communicate with the library using `mtlib` command:

```
root@brazil:/ > mtlib -l gar -qL
Library Data:
   operational state.........Automated Operational State
                             Intervention Required
   functional state..........00
   input stations............1
   output stations...........1
   input/output status.......ALL input stations empty
                             ALL output stations empty
                             Bulk input/output allowed
   machine type..............3494
   sequence number...........14724
   number of cells...........1101
   available cells...........941
   subsystems................13
   convenience capacity......10
   accessor config...........01
   accessor status...........Accessor available
                             Gripper 1 available
                             Gripper 2 available
                             Vision system operational
   comp avail status.........Primary library manager installed.
                             Primary library manager available.
                             Primary hard drive installed.
                             Primary hard drive available.
                             Secondary hard drive installed.
                             Secondary hard drive available.
                             Convenience input station installed.
                             Convenience input station available.
                             Convenience output station installed.
                             Convenience output station available.
   library facilities........00
   bulk input capacity.......50
   bulk input empty cells....Unknown
   bulk output capacity......50
   bulk output empty cells...50
   avail 3490 cleaner cycles..0
   avail 3590 cleaner cycles..25
root@brazil:/ >
```

At this point, the library can be defined to Tivoli Storage Manager. Refer to section 4.2.4, "Tivoli Storage Manager configuration" on page 118 for 3494-specific procedures for Tivoli Storage Manager.

Windows 2000 preparation

Configuring Windows 2000 servers to use the 3494 library is also very similar to what you would do for a SCSI library, again with the exception of the additional step of installing the `atldd` driver.

Configure the server as you would for communication to a SCSI library. See section "Windows 2000 preparation" on page 84 for prerequisites and instructions for a Windows Server installation. When those steps are completed, proceed to install the `atldd` driver.

Download the latest version of the `atldd` driver to a temporary directory on the server to be attached to the library. The latest version of the `atldd` driver can be obtained from:

`ftp://index.storsys.ibm.com/devdrvr/Win2000/`

The file will be named something similar to `IBMatl.4.1.5.0.exebin`. Rename the file to an appropriate `.exe` extension, such as `IBMatl.4.1.5.0.exe`. Run the file by double-clicking in Windows Explorer. The installation routine uses the standard Windows installer; follow the prompts to install the driver. The default settings should be fine.

The installation routine does not require a reboot, however, if you do not reboot, it will be necessary to start the driver manually. The "readme" file that comes with the driver describes the method to do this.

The configuration file for the driver is "`c:\winnt\ibmatl.conf`", and is a plain text file that you can edit with notepad. The format is the same as for the UNIX version, and the default file is well commented. Add a line for the library that you will be using:

```
gar    193.1.1.100    gar
```

This defines a library called "gar" at IP address 193.1.1.100, which will be known as "gar" to the host system.

Use the `mtlib` utility to verify communication with the library manager. This utility will be installed in the `%SystemRoot%\system32` directory. Typically, this will be `c:\winnt\system32`.

```
C:\WINNT\system32>mtlib -l gar -qD
Library Data:
   operational state.........Automated Operational State
                             Intervention Required
   functional state..........00
   input stations............1
   output stations...........1
   input/output status........ALL input stations empty
                             ALL output stations empty
                             Bulk input/output allowed
   machine type..............3494
   sequence number...........14724
   number of cells...........1101
   available cells...........941
   subsystems................13
   convenience capacity......10
   accessor config...........01
   accessor status...........Accessor available
                             Gripper 1 available
                             Gripper 2 available
                             Vision system operational
   comp avail status.........Primary library manager installed.
                             Primary library manager available.
                             Primary hard drive installed.
                             Primary hard drive available.
                             Secondary hard drive installed.
                             Secondary hard drive available.
                             Convenience input station installed.
                             Convenience input station available.
                             Convenience output station installed.
                             Convenience output station available.
   library facilities........00
   bulk input capacity.......50
   bulk input empty cells....Unknown
   bulk output capacity......50
   bulk output empty cells...50
   avail 3490 cleaner cycles..0
   avail 3590 cleaner cycles..25
C:\WINNT\system32>
```

At this point, communication with the library has been attained. Proceed to configure Tivoli Storage Manager in the next section.

4.2.4 Tivoli Storage Manager configuration

Once the hardware, firmware, and operating systems are prepared as described in the previous sections, Tivoli Storage Manager configuration is virtually identical to prior releases, with a couple of important exceptions.

The instructions in this section make the following assumptions:

- You have are able to “see” your library devices as described in the previous sections.
- You have loaded the drivers appropriate for your platforms, and the operating systems are configured to use those drivers.
- You have updated your hardware to a current microcode/firmware level.
- Tivoli Storage Manager is loaded on the servers, and has been updated to at least version 4.1.0.

Note that for the purposes of this publication, we used Tivoli Storage Manager version 4.1.1.

4.2.4.1 Configuring the first server

We chose to set up the RS/6000 F50 server first. Unlike SCSI library sharing, the order in which you configure the servers is not important. The library manager, in the case of the 3494, is actually a part of the 3494, and does not need to be separately defined as it does with a SCSI library.

Note

In this section, in the interest of brevity, we have elected to omit server command responses. Assume that except where noted, the server has responded with a successful completion message.

In a shared 3494 environment with shared 3590 drives, the following entries are required in the server options file, also known as `dsmserv.opt` as shown below.

```
ENABLE3590LIBRARY  YES
3494SHARED  YES
```

After inserting these entries, be sure to restart the Tivoli Storage Manager server.

On the F50 we next defined the library. As this library is to be shared with another Tivoli Storage Manager server, it was necessary to specify that different volume categories than the default should be used:

```
DEFINE LIBRARY GAR LIBTYPE=349X DEVICE=/dev/gar PRIVATECATEGORY=400
SCRATCHCATEGORY=600
```

We suggest that you refer to the publication *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409 for more detailed information on assigning volumes to volume categories in a 3494 library.

After defining the library, you should proceed to define the drives within the library. In a typical SCSI connected environment, there is usually not a problem determining which system device refers to a physical device within a library; in a SAN environment, this can be considerably more difficult. We suggest that you use the `lscfg` command to determine the serial number of each AIX “rmt” device, and compare it to the serial number physically on the drive:

.

```
(root@ost3)/=> lscfg -vl rmt4
  DEVICE            LOCATION          DESCRIPTION

  rmt4              2A-08-01          IBM 3590 Tape Drive and Medium
                                      Changer (FCP)

        Manufacturer................IBM
        Machine Type and Model......03590E1A
        Serial Number...............0000000E0019
        Device Specific.(FW)........E2DE
        Loadable Microcode Level....A0B00E33
```

In our lab scenario, we had access to a number of drives in two different 3494 libraries; however, we wished to only define this one drive to Tivoli Storage Manager. In such an instance, it is necessary to use a method such as this to determine the drive mapping for your system.

Now that we knew which drive to define to Tivoli Storage Manager, we did so, using the command:

```
DEFINE DRIVE GAR DRIVE0 DEVICE=/dev/rmt4 ONLINE=Yes
```

At this point, the library and drives are defined in the same manner that you would use in a non-shared environment. Issuing a `QUERY LIBRARY` command returns the following:

```
tsm: BRAZIL>q libr

    Library Name: GAR
    Library Type: 349X
          Device: /dev/gar
Private Category: 400
Scratch Category: 600
External Manager:
          Shared: No

tsm: BRAZIL>q drive

Library Name  Drive Name    Device Type  Device           ON LINE
------------  ------------  -----------  ---------------- -------------------
GAR           DRIVE1        3590         /dev/rmt4        Yes
```

Next, you should define the device class for the drives.

```
DEFINE DEVCLASS GARCLASS DEVTYPE=3590 FORMAT=DRIVE MOUNTLIMIT=1
MOUNTWAIT=60 MOUNTRETENTION=2 PREFIX=ADSM LIBRARY=GAR
```

Important:

A key element to remember when defining the device class for a shared 3494 environment is to explicitly set the mount limit to the number of drives that you have. When setting up the device class, the default entry for the mountlimit parameter is “drives”. For a shared Fibre Channel connected 3590 drives, this will not work correctly.

At this point, the first server’s configuration relevant to 3494 sharing is completed.

4.2.4.2 Configuring the second server

The second server was configured much as the first was, with a couple of exceptions. Most notably, it is necessary to use different volume categories when defining the library.

We began by adding the following in the server options file `dsmserv.opt`:

```
ENABLE3590LIBRARY  YES
3494SHARED  YES
```

Remember to restart the server after this change.

Next, we defined the library to the server, using the following command:

```
DEFINE LIBRARY GAR LIBTYPE=349X DEVICE=GAR PRIVATECATEGORY=112
SCRATCHCATEGORY=300
```

Note that we used different volume categories from the first server when creating the library definition.

We again refer the reader to the publication *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409 for a more thorough treatment of volume category assignment in a shared 3494 environment.

Important

Unlike shared SCSI libraries, it is not possible to share scratch volumes when sharing a 3494 library in this manner. This is due to the fact that the library manager is a part of the 3494 library, and is not resident on one of the Tivoli Storage Manager servers. When sharing is done in this manner, the Tivoli Storage Manager server databases are not synchronized.

Accordingly, if one Tivoli Storage Manager server checks out a volume from the 3494, the other server has no way to know that the volume is missing.

For this reason, it is necessary to maintain separate scratch volume pools and categories for each Tivoli Storage Manager server attached to the shared library.

As with the first server, it is important to determine which drive to define to Tivoli Storage Manager. In our lab scenario, our second Tivoli Storage Manager server was a Windows 2000 server. To ascertain which drive to define, we looked up the following key in the Windows 2000 server's registry to determine the drive's serial number, and compared this to the physical devices:

```
HKEY_LOCAL_MACHINE\HARDWARE\DEVICEMAP\Scsi\Scsi Port 3\Scsi Bus 0\Target Id
0\Logical Unit Id 2\
```

The `InquiryData` binary value has the drive serial number embedded in it. Using this, we were able to determine the drive to define to Tivoli Storage Manager.

We defined the drive using the following command:

```
DEFINE DRIVE GAR E0019 DEVICE=MT2.0.0.3 ONLINE=Yes
```

Again, it was necessary to define the device class, being careful to specify the real number of drives when defining the MOUNTLIMIT parameter of the device class, and **not** using the default value of DRIVES.

```
DEFINE DEVCLASS GARCLASS DEVTYPE=3590 FORMAT=DRIVE MOUNTLIMIT=1
MOUNTWAIT=60 MOUNTRETENTION=60 PREFIX=ADSM LIBRARY=GAR
```

This completes the IBM 3494 sharing-specific setup instructions.

4.2.5 Operating a shared IBM 3494 tape library

In the 3494 shared environment as we had set it up, operations were largely the same as you would see with a single attached server. As each server could only see and access tapes assigned to it, tape operations behave as if the library were dedicated to that server.

4.2.5.1 Drive contention in a shared 3494

By defining only one drive in our shared IBM 3494 scenario, it was very easy to create situations where drive contention arose.

As the two Tivoli Storage Manager servers act independently, and no synchronization is done between them, there is no pre-emption for higher priority processes.

When one server obtains the use of a drive, other servers requiring the use of that drive will find that the drive is locked, and begin to poll the drive. The activity log output from the second server shown below illustrates this behavior:

```
09/12/2000 16:27:30      ANR8848W Drive E0019 of library GAR is inaccessible;
server has begun polling drive.
```

When the drive becomes free again, the following will be entered into the activity log:

```
09/12/2000 16:28:01      ANR8839W Drive E0019 of library GAR is accessible.
```

The pending operation will then proceed on the second server.

Note however, that this behavior is governed by the MOUNTWAIT parameter; if this value is set too low, the pending transaction will time out before the server with ownership of the drive releases it. The MOUNTWAIT parameter is changed by updating the device class like this:

```
UPDATE DEVCLASS GARCLASS MOUNTWAIT=120
```

Some experimentation may be required for your particular environment.

4.2.5.2 Monitoring drive activity

In the event of a hardware failure on a drive which is configured in this fashion, the Tivoli Storage Manager server which is either using the drive, or attempts to acquire the drive first will flag the drive as unavailable. However, since the Tivoli Storage Manager database is not shared between servers, the new status information is not propagated to the other servers.

Due to this, if a drive in a 3494 library fails, the status of the drives as known to the Tivoli Storage Manager servers should be checked, and updated if they do not accurately reflect the status of the drive.

The server option that affects how many times a Tivoli Storage Manager server attempts to acquire a drive is DRIVEACQUIRERETRY. The default for this parameter is FOREVER. Changing this parameter to a different value, however, should be done with caution; if it is set to a low value, operations may fail if they are unable to obtain the drive before the retry count expires.

For example, assume that server A is performing a migration operation to a 3590 Fibre Channel drive in a 3494 library, and that the drive fails. The Tivoli Storage Manager server on server A will flag the drive as unavailable. When server B attempts to gain access to the drive, it will be unaware that the drive is not ready for operation, and will begin to poll the drive. Assuming that DRIVEACQUIRERETRY is set to the default of “forever”, server B will never set the drive status to unavailable. Server B will continue to attempt to retry the operation until the MOUNTWAIT setting in the device class for the drives is exceeded.

Chapter 5. LAN-free client data transfer

This chapter discusses LAN-free client data transfer solutions using Tivoli Storage Manager and Storage Agent. In particular, we will cover:

- Main concepts of LAN-free client data transfer using the new Tivoli Storage Agent available since TIvoli Storage Manager Version 4.1
- Setup and configuration of TDP for Exchange and TDP for R/3 to use LAN-free client data transfer
- Operations and considerations when using client data transfer

5.1 Concepts of LAN-free client data transfer

In this section we explain the main concepts of Tivoli Storage Manager LAN-free client data transfer.

5.1.1 LAN based data transfer

In most of today's environments a LAN-based data transfer (see Figure 43) is being used. This model is based on one (or more) Tivoli Storage Manager servers, which utilize a local SCSI connection to a storage device (disk arrays; tape and optical libraries).

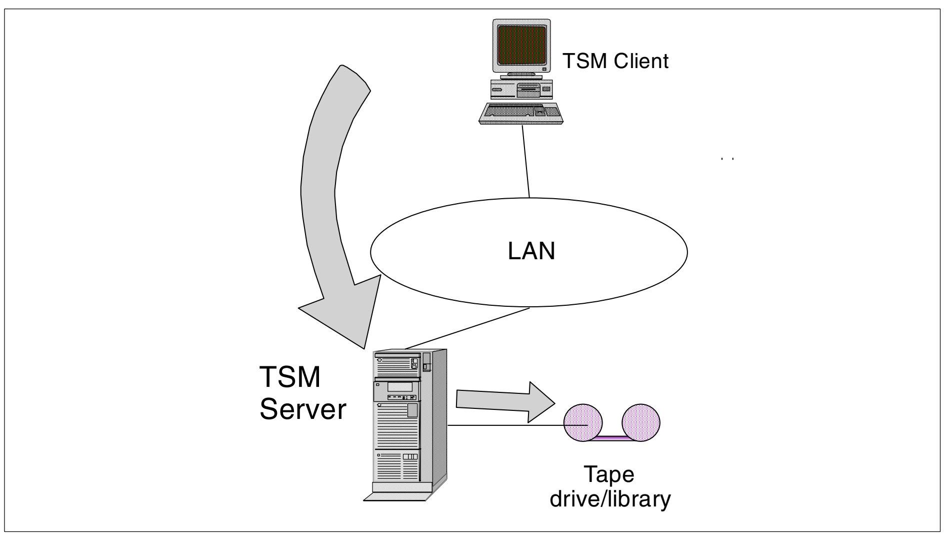

Figure 43. Common LAN-based backup/restore scenario

The Tivoli Storage Manager server is connected to a LAN. When a Tivoli Storage Manager client wants to backup/restore its data to a Tivoli Storage Manager server, it opens a session via LAN to Tivoli Storage Manager server and begins to send data.

In most cases, a LAN is *not* designed for the huge data transfers needed to backup/restore business critical data.

Note

You can get more information about SAN concepts in Chapter 1, “SAN overview” on page 1 and in Chapter 3, “Building a SAN environment” on page 51. We recommend that you also read the following IBM Redbooks:

- *Introduction to Storage Area Networks*, SG24-5470
- *Tape future in fabrics,* SG24-5474

5.1.2 LAN-free client data transfer

A SAN is designed to eliminate this bottleneck of LANs. It is a dedicated high-speed network based on Fibre Channel protocol.

In Figure 44 you can see how the infrastructure has changed for the LAN-free scenario.

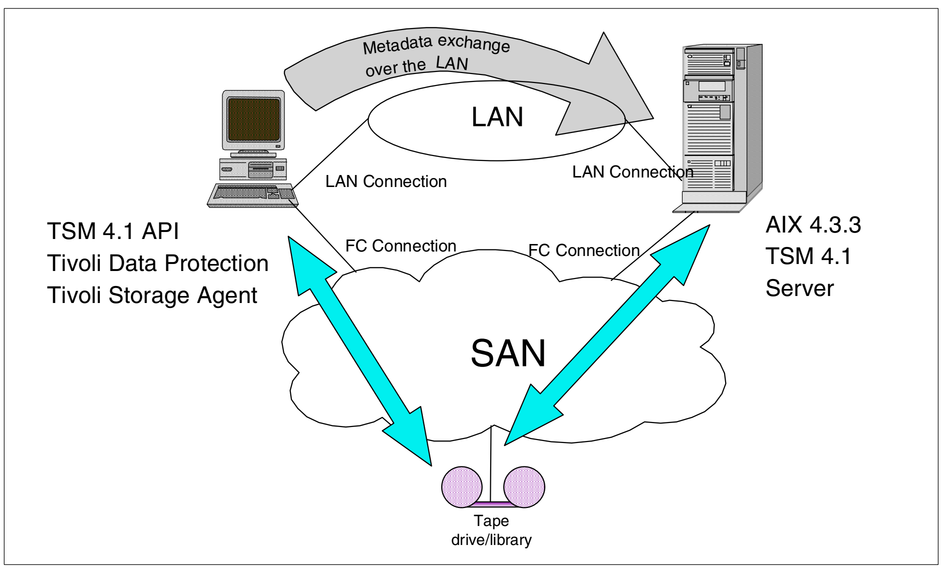

Figure 44. TSM LAN-free storage solution

We still need a LAN for exchanging control information between Tivoli Storage Manager client and Tivoli Storage Manager server. But the raw data will be dispatched on the client to the SAN and transferred across the SAN directly to the library. The term, *control information*, means that the client needs to access the Tivoli Storage Manager server's database.

To be able to send the data from the Tivoli Storage Manager client across the SAN to the library, we need some piece of software on the client machine, which basically will act as an "interface" between the client machine, the SAN, and the Tivoli Storage Manager server. In the Tivoli Storage Manager environment, the software that does this is called the *Storage Agent*, and we will explain its function more closely in the next section.

5.1.3 The Storage Agent

To achieve LAN-Free data transfer, a new feature has been introduced in the Tivoli Storage Manager family — Tivoli Storage Manager / Managed Systems for SAN (the Storage Agent).

The Storage Agent software runs in the following environments:

- Windows NT4.0, service pack 6 and higher
- Windows 2000 Server, build 2195 and higher

In our scenario, we had our Tivoli Storage Manager server running on AIX, but you can use any server platform since release 4.1 and higher.

The client machine is running a Storage Agent, which basically enables the LAN-free data transfer. During testing in our lab, there were only two products supported with a Storage Agent:

- Tivoli Data Protection (TDP) for MS Exchange
- Tivoli Data Protection (TDP) for R/3

In the future, other TDP clients will be supported also.

The Storage Agent works in the following way. When Tivoli Data Protection invokes a backup:

1. A named pipe session (on Windows platform) is opened to the Storage Agent; the name of the named pipe is hardcoded in TDP code. The only communication protocol supported is named pipe.
2. The Storage Agent opens a server-to-server connection to Tivoli Storage Manager server via the LAN; the only supported communication protocol at this time is TCP/IP.

3. TDP sends a begin-transaction verb to Storage Agent. When the transaction is successfully committed, Storage Agent opens a SAN data transfer to a tape storage device. Otherwise, if the first transaction fails, the communication will be switched over to the LAN, and the SAN data transfer will not be used.
4. The LAN server-to-server communication between the Storage Agent and Tivoli Storage Manager server is used by the Storage Agent for accessing Tivoli Storage Manager server database. That is because the Storage Agent does not have its own database, nor a recovery log. You can imagine the Storage Agent as a “truncated” Tivoli Storage Manager Server.

When restoring data, the restore session is always opened through a Storage Agent. If any data was backed up through the SAN, the restore will also go through the SAN. If any data was backed up through a LAN, then the LAN connection will be used during restore. That means, if you have backed up some data using the SAN, you can still access this data through the LAN, if needed.

Note that TDP opens only one session to the Storage Agent. Multiple concurrent sessions are not yet supported. The destination storage pool always has to be a sequential access storage pool, because the Storage Agent is not able to communicate to disk devices over the SAN.

The plan also includes making the Storage Agent available for UNIX platforms. This concept will be similar to implementations on Windows systems, with some changes. For example, a shared memory protocol will be used to open a session from TDP to the Storage Agent, instead of the named pipe communication protocol, which is not supported on UNIX.

5.2 Setting up LAN-free client data transfer

Before setting up LAN-free data transfer, you should have a Tivoli Storage Manager 4.1 server configured and running, since it must be able to communicate with any tape device via the SAN. You can have any server platform running; the only requirement is that it must be Release 4.1 or later.

For information on how to configure your SAN environment, as well as how to configure HBAs on different platforms, read Chapter 3, “Building a SAN environment” on page 51.

We assume that you already have installed the application that you plan to back up on the client computer (MS Exchange, SAP R/3).

5.2.1 Preparing TSM server for LAN-free configuration

Basically, you need to have a Tivoli Storage Manager 4.1 server properly configured in order to be able to use your SAN attached library. For more information how to configure Tivoli Storage Manager server for using SAN attached libraries, see Chapter 4, “Tape library sharing” on page 75.

A copy group definition in the appropriate management class, which you will be using for LAN-free backups always points to a sequential access storage pool. You cannot use disks for LAN-free backup using a Storage Agent.

You also need to have a client node registered in order to be able to connect a TDP client to Tivoli Storage Manager server.

In our environment, we ran Tivoli Storage Manager server 4.1.1.0. The machine ran AIX 4.3.3., installed maintenance level 4.

To prepare your Tivoli Storage Manager server for LAN-free configuration, you need to do the following:

- Define library and drives. The library must be defined as a SHARED library, otherwise the Storage Agent will not be able to communicate with it. Enter the following commands at the Tivoli Storage Manager server prompt to define the library and drives:

```
DEF LIBR 3570LIB LIBT=SCSI DEVICE=/dev/rmt0.smc SHARED=YES
DEF DR 3570LIB DRIVE0 DEVICE=/dev/rmt0
DEF DR 3570LIB DRIVE1 DEVICE=/dev/rmt1
```

  If you have already a library defined on your server, make sure that the parameter SHARED is set to YES. Note that you cannot use the `update library` command to update this parameter. You must remove the drives definition, delete the library, and then define it again with the SHARED parameter set to YES.

- Define device class; in our example, we are defining a device class for the 3570library:

```
DEF DEVCLASS 3570CLASS DEVTYPE=3570 MOUNTRET=2 MOUNTL=2 LIBRARY=3570LIB
```

We set the MOUNTRETENTION to a very low value (2 minutes) and the MOUNTLIMIT parameter according to the number of physical drives in the library. These values seem to us to be optimal for the library sharing operation. For a more detailed explanation, read Section 4.1.7, “Using a shared SCSI library with Tivoli Storage Manager” on page 103.

- Define a storage pool in the device class you created previously.

```
DEF STG TAPEPOOL 3570CLASS
```

You also need to have labeled volumes checked in your library, a defined policy, and a client node registered to the appropriate policy domain. If you need further explanation on how to do this, read the documentation supplied to you with the Tivoli Storage Manager server.

When we registered the client node TDPDIOMEDE to our server, we set the parameter COMPRESS to NO to prevent TDP from compressing the data on the client.

5.2.2 Verify configuration of SAN devices on the client

We have tested LAN-free transfer both on Windows NT and Windows 2000 server. In this section we will show you what device configuration you need to check before you can install the Storage Agent and TDP.

5.2.2.1 Windows NT checklist for SAN configuration

1. Check your operating system to see if it has the latest service pack installed; if not, install it.
2. Open the Control Panel, SCSI devices, and check if the operating system sees your Fibre Channel HBAs, as shown in Figure 45.

   In our test environment we ran Windows NT 4.0, service pack 6. To connect this machine to a SAN, we used Qlogic HBA adapter 2100A series.

   Make sure you always use latest driver code. For the latest Qlogic Fibre Channel adapter device drivers, look at the following Web site:

   http://www.qlogic.com/bbs-html/csg_web/adapter_pages/fc_adapters.html

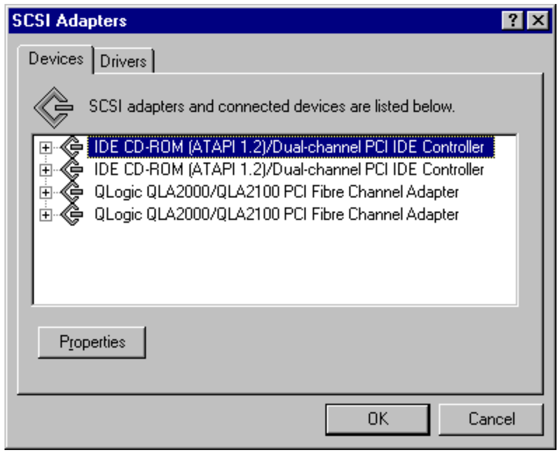

Figure 45. Installing Qlogic FC adapter drivers

3. Check if the adapter is started and is able to see devices on the SAN as shown in Figure 46.

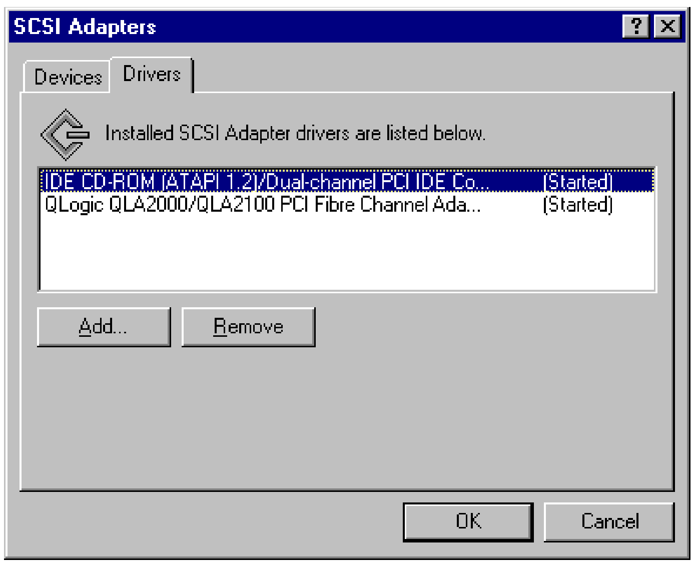

Figure 46. Qlogic device driver configuration

In Figure 47 you can see that the Fibre Channel adapter is able to see SAN-connected devices.

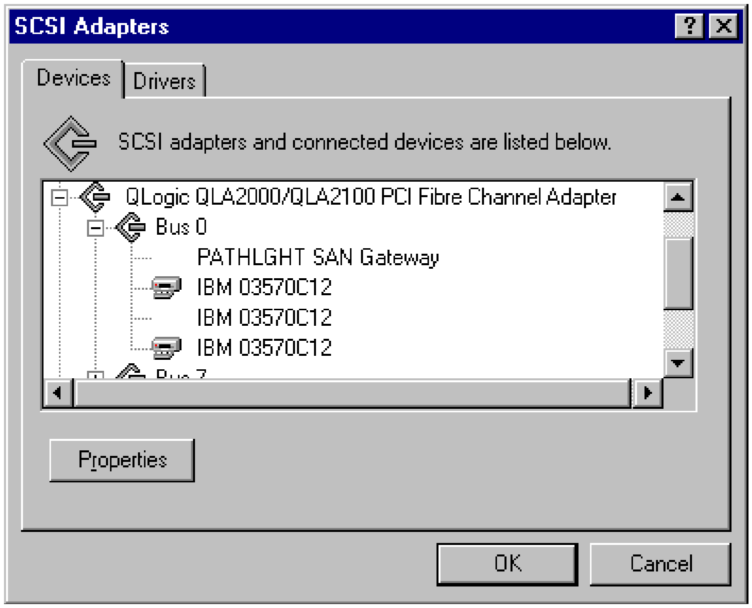

Figure 47. Qlogic device driver configuration, devices visible to HBA

5.3 Installing and configuring the Storage Agent

You need to go through these steps to configure the Storage Agent:

- Install the Storage Agent Software.
- Modify `DSMSTA.OPT` file.
- Set up server-to-server communication.
- Obtain device information from `TSMDLST.EXE`.
- Define drive mapping on Tivoli Storage Manager server.
- Install the Storage Agent as a service.

Keep in mind that you cannot install the Storage Agent on the same machine where a Tivoli Storage Manager server resides. In this case, you will not be able to install the Storage Agent, as shown in Figure 48.

Figure 48. Storage Agent not on the same machine as a TSM server

We have also tried to stop the Tivoli Storage Manager server and then run the Storage Agent installation routine, but we got the same error as shown in Figure 48. We had to uninstall our Tivoli Storage Manager server software; then we could install the Storage Agent.

In our opinion, it does not make sense to install both Tivoli Storage Manager server and the Storage Agent on one machine. This is because when you run the Tivoli Storage Manager server on a SAN connected machine, you can configure it as a library client in order to be able to do LAN-free backups from this machine. The Storage Agent is designated to make LAN-free data possible on your client machines, and that is most probably the reason you cannot install it on the same machine where Tivoli Storage Manager server already resides.

5.3.1 Installing the Storage Agent Software

To install the Storage Agent software, follow these steps:

1. From Windows Explorer, click on **setup.exe**. Just as with other Windows-based products, a setup wizard will guide you through the installation process. If you do not to want install the Storage Agent to the default destination, change the destination folder (Figure 49).

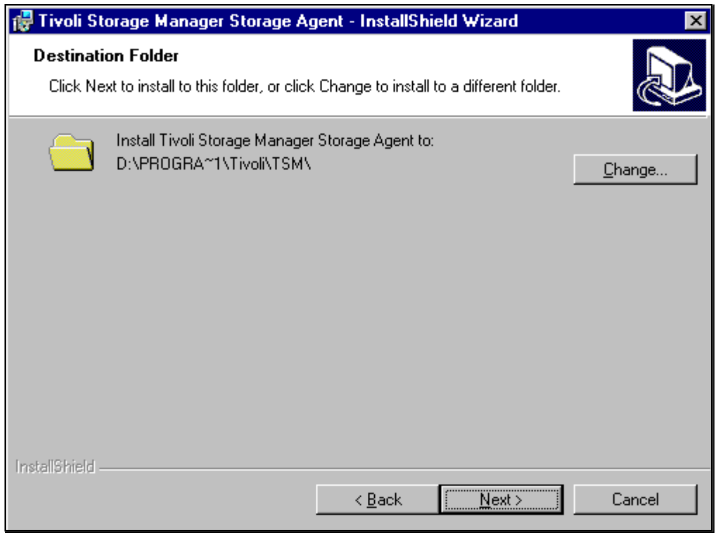

Figure 49. Installing the Storage Agent

2. Select custom installation (Figure 50).

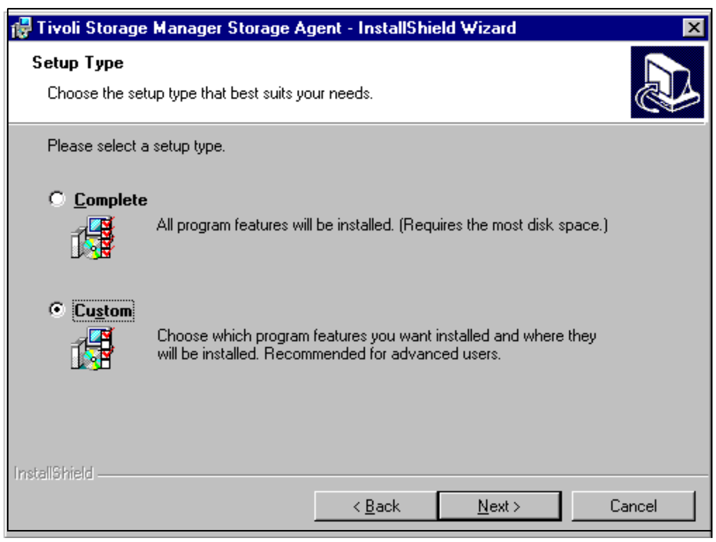

Figure 50. Installing the Storage Agent, choosing setup type

3. Select the components you want to install (Figure 51).

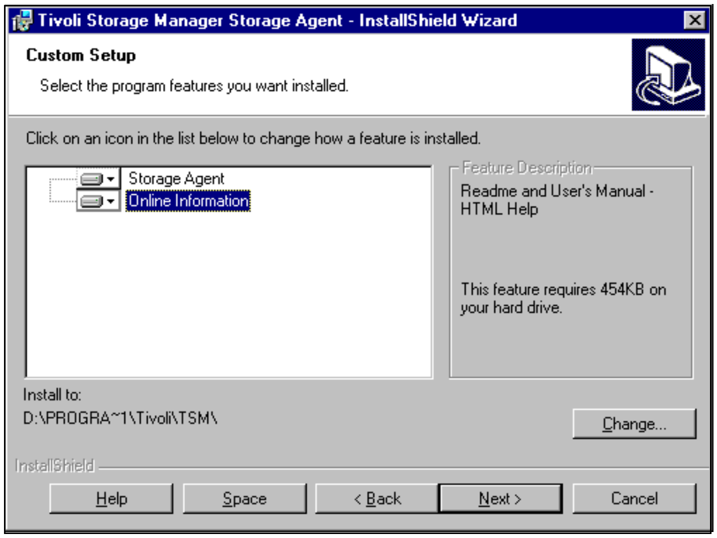

Figure 51. Installing the Storage Agent, selecting components

4. Click the **Install** button to begin installation (Figure 52).

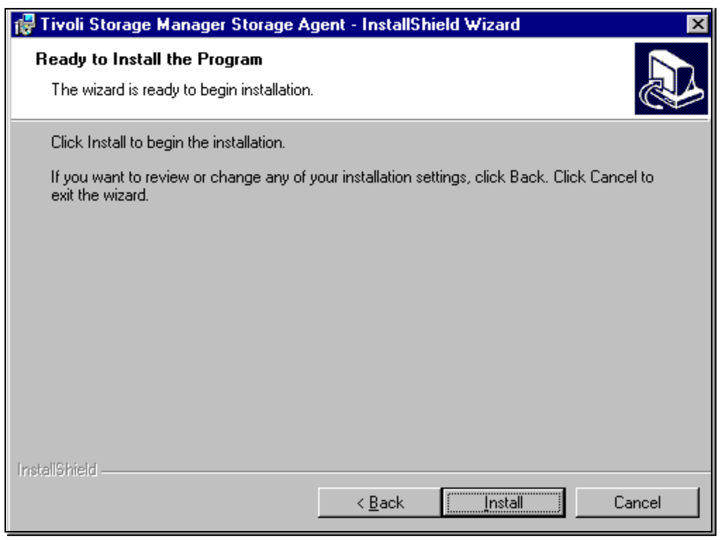

Figure 52. Installing the Storage Agent, copying files

5. A warning message may appear, if there is no HTML help installed on your system. Click **OK** to install HTML help. If you do not install HTML help, you will not be able to read the documentation provided with the Storage Agent (Figure 53).

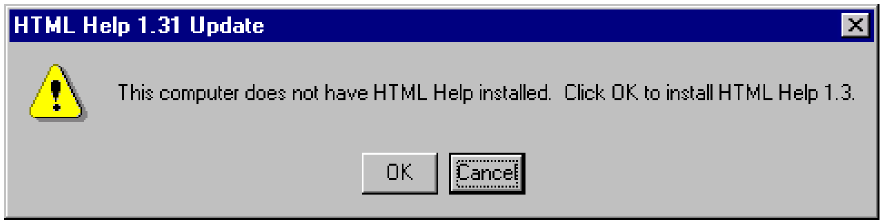

Figure 53. Installing the Storage Agent, install HTML help

6. Click **Finish** to complete the installation (Figure 54).

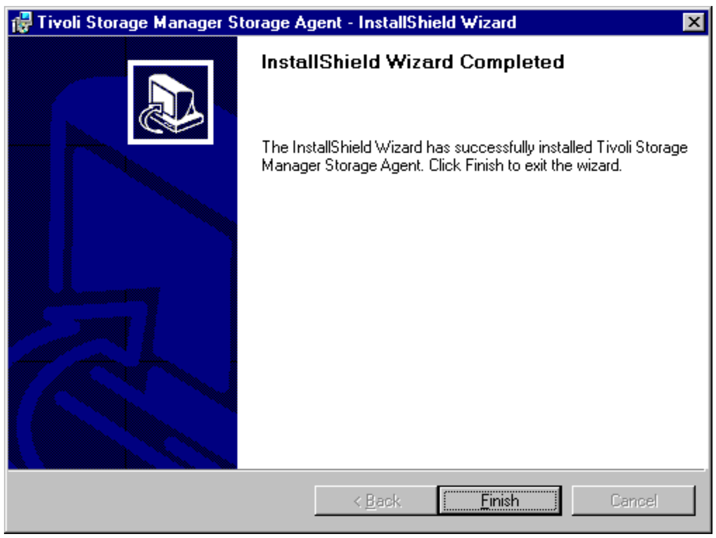

Figure 54. Installing the Storage Agent, finishing installation

7. A reboot is required after installing the Storage Agent. Close all your applications and click **Yes** to reboot your computer. Note that when you click **Yes**, your computer starts rebooting immediately (Figure 55).

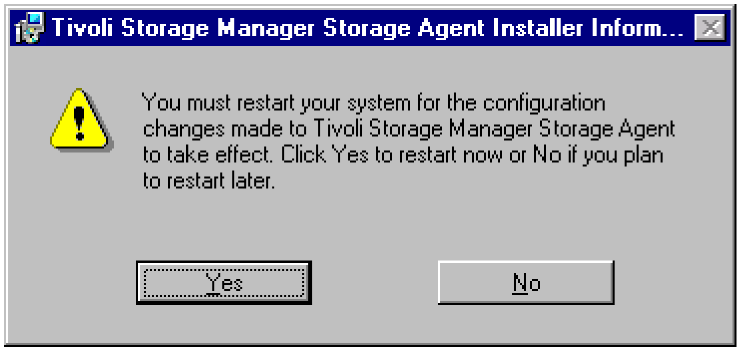

Figure 55. Installing the Storage Agent, rebooting computer

During our tests we used the Storage Agent Version 4.1.1.0.

You can download the latest Storage Agent code from:

`ftp://index.storsys.ibm.com/tivoli-storage-management/maintenance/server/v4r1/NT/LATEST/`

5.3.2 Modifying DSMSTA.OPT file

Open and edit the `DSMSTA.OPT` file. You need to enable the named pipe and TCP/IP communication protocols. The named pipe protocol is used by TDP to communicate with the Storage Agent, and TCP/IP must be used for communication between the Storage Agent and Tivoli Storage Manager server.

Here is an example of how we have modified the `DSMSTA.OPT` file:

```
COMMmethod TCPIP
COMMmethod NAMEDPIPE
DEVCONFIG devconfig.txt
ENABLEALLCLIENTS yes
```

There is an undocumented option `enableallclients` which enables other clients to talk to the Storage Agent. Note that at the time of writing, the Storage Agent only supports TDP for MS Exchange and TDP for SAP R/3. You should **not** use the Storage Agent with any unsupported products.

5.3.3 Start AdsmScsi device driver

Now you must make sure that the AdsmScsi service is running. In case it is not running, start the service. If the service cannot be started, there may be a problem with your HBA configuration.

In this case, make sure that the HBA device driver is loaded and that the adapter is able to see your devices.

To start the `AdsmScsi` service, issue the following command from the Storage Agent directory:

```
net start AdsmScsi
```

5.3.4 Define drive mapping

Once your AdsmScsi service is running, run the `TSMDLST` command from the Storage Agent folder. You should see output similar to the following:

```
D:\Program Files\Tivoli\TSM\storageagent>tsmdlst

Computer Name:      DIOMEDE
TSM Device Driver:  Running

TSM Device Name  ID   LUN  Bus  Port   TSM Device Type  Device Identifier
-------------------------------------------------------------------------
mt0.2.0.2        0    2    0    2      3570             IBM     03570C12
lb0.3.0.2        0    3    0    2      LIBRARY          IBM     03570C12
mt0.4.0.2        0    4    0    2      3570             IBM     03570C12
```

Now you must define drive mapping on Tivoli Storage Manager server side. This is needed by the Storage Agent to be able to communicate with the drive. The syntax of the command follows:

```
DEF DRIVEMAPPING <STORAGE_AGENT_NAME> <LIBRARY_NAME> <DRIVE_NAME> <DEVICE>
```

The next screen shows the sample syntax of the `DEFINE DRIVEMAPPING` command in our scenario:

```
DEF DRIVEM storagent 3570lib drive0 device=mt0.2.0.2
DEF DRIVEM storagent 3570lib drive1 device=mt0.4.0.2
```

If you do not execute the `DEFINE DRIVEMAPPING` command on the Tivoli Storage Manager server, the Storage Agent will not be able to access drives in the library. In this case, TDP will open a regular TCP-IP session to your Tivoli Storage Manager server, and LAN data transfer will be used.

> **Important**
>
> The device parameter you must specify when issuing the `DEFINE DRIVEMAPPING` command is the device name which you got from the `TSMDLST` command. Do **not** use Tivoli Storage Manager server's devices.

5.3.5 Setting up server-to-server communication

The Storage Agent needs to exchange metadata with Tivoli Storage Manager server. For this purpose it uses server-to-server communication.

On Tivoli Storage Manager server, define parameters for server-to-server communication. Next screen shows a set of commands we have used in our environment. You should change these parameters according to your environment:

```
SET SERVERNAME BRAZIL
SET SERVERPASSWORD BRAZIL
SET SERVERHLADDRESS 193.1.1.11
SET SERVERLLADDRESS 1500
DEFINE SERVER SA SERVERPASSWORD=ADMIN HLA=193.1.1.16 LLA=1500 COMM=TCPIP
```

On the client machine, go to the Storage Agent directory and issue:

```
DSMSTA SETSTORAGESERVER MYNAME=SA MYPASSWORD=admin SERVERNAME=BRAZIL
SERVERPASSWORD=BRAZIL HLADDRESS=193.1.1.11 LLADDRESS=1500
```

After this command completes, you can look into the Tivoli Storage Manager Server's activity log, where you can find, that the Storage Agent has started and ended a session to the Tivoli Storage Manager Server.

5.3.6 Installing the Storage Agent as a service

Now install the Storage Agent as a service. To do this, run the following command:

```
c:\program files\tivoli\tsm\storageagnet\install "TSM Storage Agent"
          "d:\program files\tivoli\tsm\storageagent\dstasvc.exe"
```

Make sure you have entered the correct path to your Storage Agent folder. Otherwise you will not be able to start the Storage Agent service.

If you plan to run the Storage Agent under a specific Windows NT account, you can modify the command as follows:

```
c:\program files\tivoli\tsm\storageagnet\install "TSM Storage Agent"
"c:\program files\tivoli\tsm\storageagent\dstasvc.exe" <accountname>
<password>
```

Note that you must set up this specific account before you can start the Storage Agent service.

Once you have started the Storage Agent, issue the `QUERY SESSION` command on the Tivoli Storage Manager server. You can see that the Storage Agent has opened four TCP-IP sessions to the Tivoli Storage Manager server.

```
tsm: BRAZIL>q ses

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- ---------------
     1 ShMem  Run      0 S    3.6 K     187 Admin AIX      ADMIN
     2 ShMem  IdleW    8 S    5.4 K     425 Admin AIX      ADMIN
     3 Tcp/Ip IdleW   13 S    3.6 K   3.4 K Serv- Windows  SA
                                              er    NT
     4 Tcp/Ip IdleW   12 S    1.0 K     723 Serv- Windows  SA
                                              er    NT
     5 Tcp/Ip IdleW   12 S    8.8 K   7.4 K Serv- Windows  SA
                                              er    NT
     6 Tcp/Ip IdleW   12 S      166     527 Serv- Windows  SA
                                              er    NT
```

5.3.7 Connecting to the Storage Agent

You can use the administrative client to connect to the Storage Agent. The user ID and password are the same as on your Tivoli Storage Manager server. Here, you can issue commands like `QUERY SESSION`, `QUERY PROCCES`, `QUERY ACTLOG`.

5.3.8 Using sample API client code with the Storage Agent

You can compile the sample API client code to verify that you can connect to the Storage Agent. Do not forget to specify the option ENABLELANFREE option in its `DSM.OPT` file.

For information on how to build the sample application, read the product manual *Using the Application Program Interface*, SH26-4081.

5.4 MS Exchange LAN-free backup and recovery

In this section we will describe how to set up the LAN-free client data transfer for MS Exchange. As a prerequisite, you must have an MS Exchange server environment up and running.

Configuring LAN-free data transfer on the client machine consists of the following steps:

- Installing the Version 4.1 client API
- Installing and configuring TDP
- Installing and configuring the Storage Agent

5.4.1 Installing TSM client API

To be able to setup LAN-free data transfer, you need to use 4.1 client API code. Download the latest code from:

`ftp://index.storsys.ibm.com/tivoli-storage-management/maintenance/client/v4r1/Windows/i386/LATEST/`

During our tests, we used PTF IP22088.

Note that for LAN-free data transfer, you only need to install the API runtime. You do not need to install the backup/archive client and administrative client, but you can use the administrative client to connect to the Storage Agent.

There is no need to modify anything in the configuration file `DSM.OPT` for the API. The API only has to be installed on your machine.

5.4.2 Installing and configuring TDP for MS Exchange

To install TDP for MS Exchange, you need to go through the following steps:

1. From Windows Explorer, click on the **setup.exe**. Just like by other Widows-based products, a setup wizard will guide you through the installation process (Figure 56).

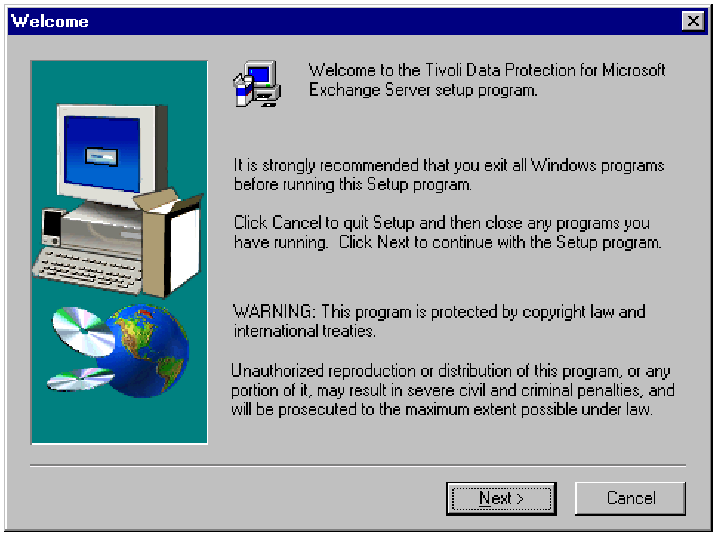

Figure 56. Installing TDP for MS Exchange

2. Choose the destination where you want to install TDP (Figure 57).

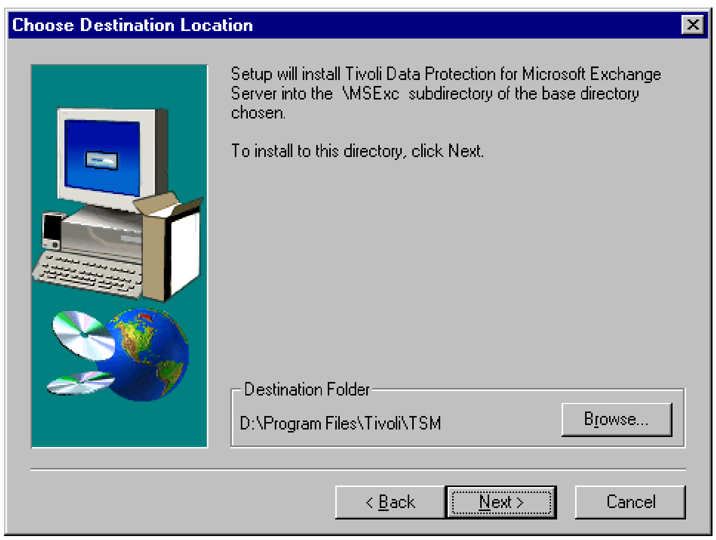

Figure 57. Installing TDP for MS Exchange, choose installation folder

3. Select **Custom** installation (Figure 58).

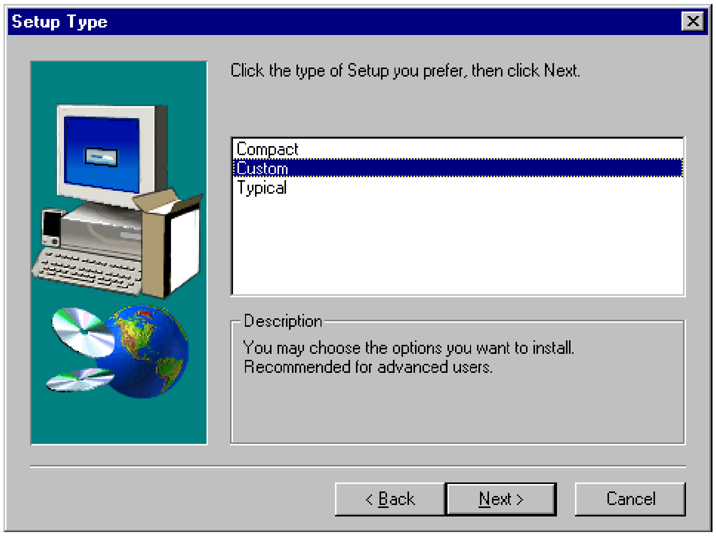

Figure 58. Installing TDP for MS Exchange, setup type

4. Select the components which you want to install (Figure 59).

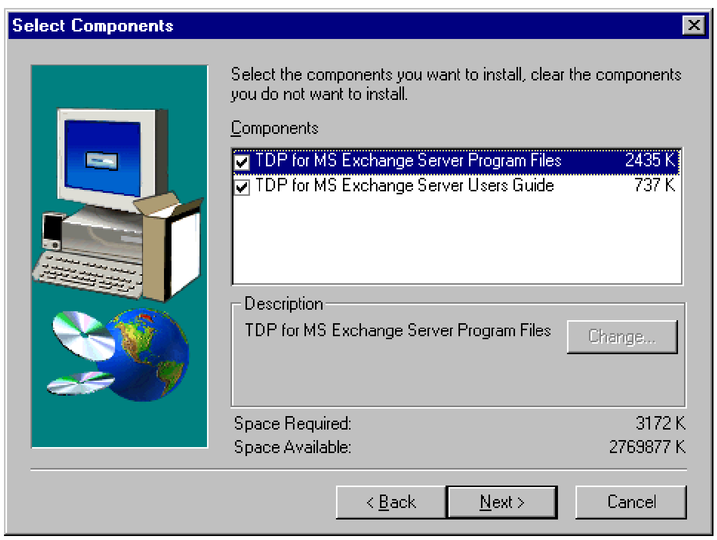

Figure 59. Installing TDP for MS Exchange, selecting components

5. Click **Finish** to complete the installation (Figure 60).

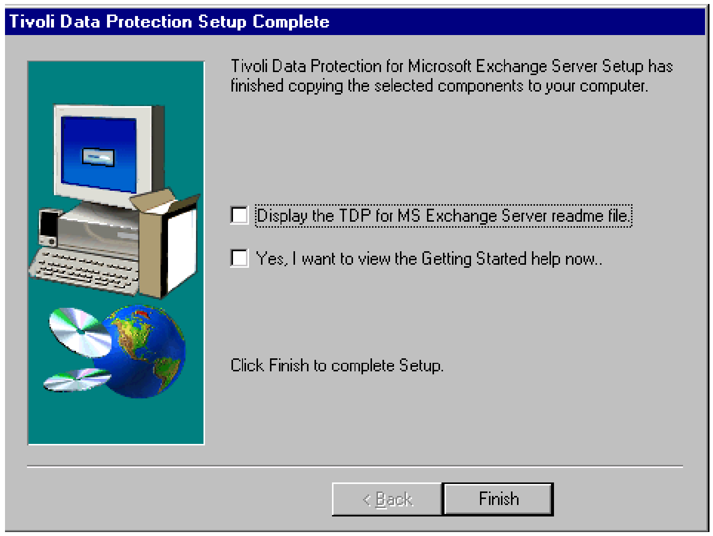

Figure 60. Installing TDP for MS Exchange, finishing installation

It is recommended that you reboot you computer after installing TDP. After reboot, open and edit the TDP's DSM.OPT file. You need to specify the following options when setting up TDP for LAN-free data transfer:

NODENAME	TDPDIOMEDE
COMMMETHOD	TCPIP
TCPSERVER	BRAZIL
TCPPORT	1500
ENABLELANFREE	YES

Note that you cannot set the ENABLELANFREE option from the TDP's graphic user interface; you need to modify the DSM.OPT file directly.

For our tests we used PTF 21909 of TDP for MS Exchange. You can download the latest code from the following FTP archive:

`ftp://index.storsys.ibm.com/tivoli-storage-management/maintenance/tivoli-data-protection/ntexch/`

5.4.3 Backing up/restoring MS Exchange using LAN-free method

For end users or operators who take backups with the TDP interface, there is no difference in how to use the TDP in SAN and non-SAN environments. In the following sections, we describe how to use TDP for backup and restore, as well as how to find out that the data transfer is going through the SAN.

In our book, we do not cover different methods of backing up, restoring, or protecting the MS Exchange server. If you need such information, please read the IBM Redbook, *Using ADSM to Back Up and Recover Microsoft Exchange Server*, SG24-5266.

5.4.3.1 Backing up MS Exchange

Start the TDP client from the Start menu on your computer. A screen similar to Figure 61 will appear:

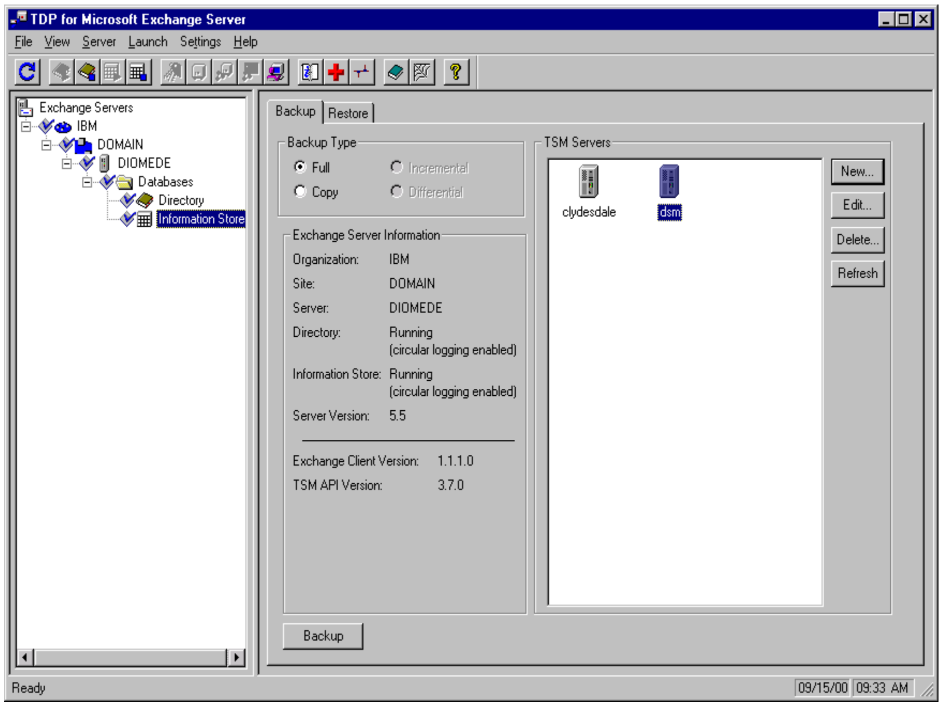

Figure 61. Backing up MS Exchange using TDP

Select **Directory** and **Information Store**, and click on the backup icon. TDP will now open a named pipe session to the Storage Agent. You can connect to the Storage Agent with the command line administrative client.

The next screen shows the output of the `QUERY SESSION` command we issued in the Storage Agent. You can see that the TDP session number 14 is started, and that it sends the data to the Storage Agent. The Storage Agent transfers the data directly to the tape drive through the SAN. You can also see that there are four sessions opened for the Storage Agent (session numbers 3,4, 8 and 15) which are used for accessing Tivoli Storage Manager server's database and library sharing:

```
tsm: SA>q ses

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- ---------------
     3 Tcp/Ip Start    0 S   27.7 K  29.6 K Serv-          SA
                                             er
     4 Tcp/Ip IdleW    0 S   18.6 K     166 Serv-          SA
                                             er
     5 Tcp/Ip Run      0 S  148.4 K   2.8 K Admin WinNT    ADMIN
     8 Tcp/Ip Start    0 S   28.0 K  27.2 K Serv-          SA
                                             er
    14 Named  SendW    0 S  378.8 M     234 Node  TDP      TDPDIOMEDE
        Pipe                                       MSExchg
                                                   NT
    15 Tcp/Ip Start    0 S    8.5 K   8.5 K Serv-          SA
                                             er
```

If you issue the `QUERY SESSION` command on the Tivoli Storage Manager server — as shown in the next screen — you can see that a TDP session is also started (session number 41 for node TDPDIOMEDE), but only a very small amount of data is being received/transferred. That is because the data is transferred through the SAN and not through the LAN (TCP/IP), and this session is opened only for the metadata exchange.

```
tsm: BRAZIL>q ses

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
     3 ShMem  Run      0 S   52.3 K      59 Admin AIX      ADMIN
    26 ShMem  Run      0 S  330.5 K  10.4 K Admin AIX      ADMIN
    33 Tcp/Ip IdleW  1.3 M   29.6 K  27.7 K Serv- Windows  SA
                                             er    NT
    34 Tcp/Ip IdleW  1.3 M      166  18.6 K Serv- Windows  SA
                                             er    NT
    36 Tcp/Ip IdleW  4.0 M   27.2 K  28.0 K Serv- Windows  SA
                                             er    NT
    41 Tcp/Ip IdleW  4.0 M    6.8 K   2.0 K Node  TDP      TDPDIOMEDE
                                                   MSExchg
                                                   NT
    42 Tcp/Ip IdleW  4.0 M    8.5 K   8.5 K Serv- Windows  SA
                                             er    NT
```

5.4.3.2 Restoring MS Exchange

Start the TDP client from the Start menu on your computer, after TDP's graphical user interface appears, and click on **Restore** (Figure 62):

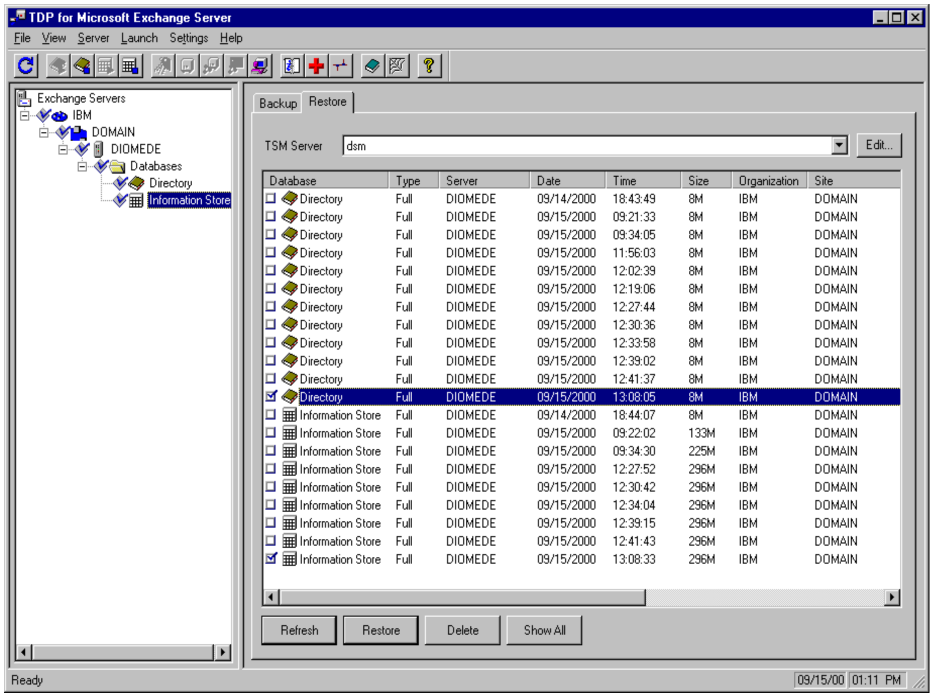

Figure 62. TDP for MS Exchange — select pair Directory and IS to restore

In our test environment we only took full backups of our Exchange Server. For restore, we choose the latest backup of **Directory** and **Information Store**.

Before restore can be started, TDP has to stop the Exchange Directory Service (Figure 63).

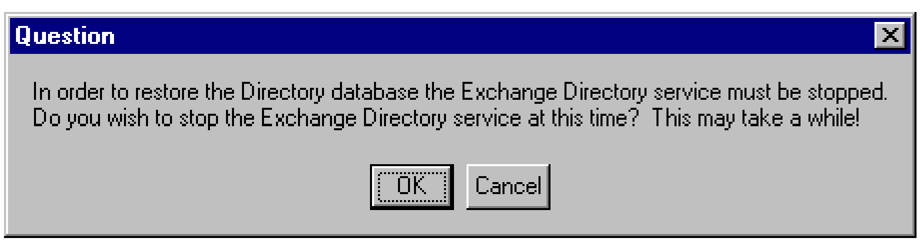

Figure 63. TDP for Exchange — stop the Exchange Directory service

Now you must select restore options for full or copy backups as shown in Figure 64.

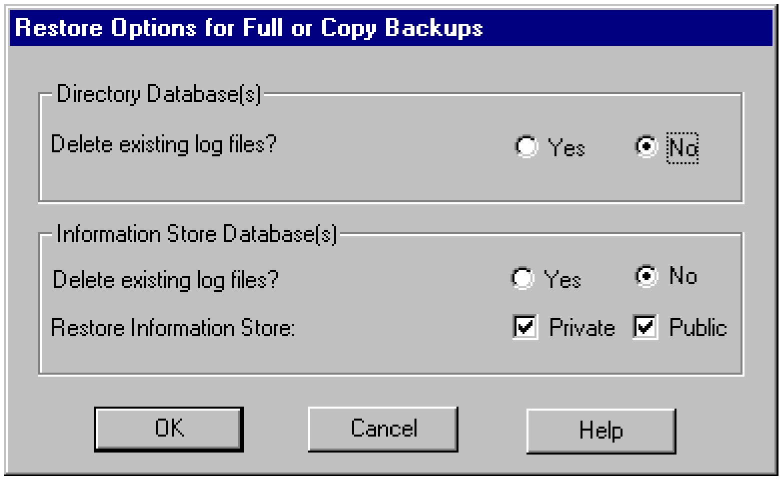

Figure 64. TDP for Exchange, selecting restore options

After you click the **OK** button, a session to the Storage Agent will be opened, and the TDP will begin to send data. Again, you can use the Tivoli Storage Manager administrative command line interface to connect to the Storage Agent and issue the `QUERY SESSION` command. In the *bytes received* column you can see how much data has been transferred from the library to the TDP.

```
tsm: SA>q ses

  Sess  Comm.   Sess     Wait    Bytes    Bytes  Sess   Platform  Client Name
Number  Method  State    Time     Sent     Recvd  Type
------  ------  ------  ------  -------  -------  -----  --------  ------------
     1  Tcp/Ip  Start     0 S    9.0 K   11.4 K  Serv-            SA
                                                  er
     3  Tcp/Ip  Start     0 S   41.3 K   46.5 K  Serv-            SA
                                                  er
     4  Tcp/Ip  IdleW     0 S   18.9 K      166  Serv-            SA
                                                  er
     7  Tcp/Ip  Start     0 S   79.1 K   78.1 K  Serv-            SA
                                                  er
    11  Tcp/Ip  Run       0 S  122.0 K    2.4 K  Admin  WinNT     ADMIN
    31  Named   Run       0 S   22.5 M      234  Node   TDP       TDPDIOMEDE
         Pipe                                            MSExchg
                                                         NT
    32  Tcp/Ip  Start     0 S    7.9 K    7.8 K  Serv-            SA
                                                  er
```

If you need to restore your data through LAN, you must stop the Storage Agent before you start the TDP client. When the SAN connection fails during restore, the behavior is the same as described in Section 5.4.3.1, “Backing up MS Exchange” on page 145.

5.5 LAN-free backup and recovery of R/3

In this section we will describe how to set up and use LAN-free data transfer for SAP R/3 using TDP for R/3. At the time of writing, the only supported platform for LAN-free data transfer is Windows NT, so we have only tested R/3 on this platform.

5.5.1 Preparing your environment for LAN-free method

Before you start to install and configure the TDP for R/3, make sure that you can access all of the SAN devices you plan to use. Also make sure that the client API Version 4.1 and the Storage Agent are installed and configured properly.

For information on how to install and set up your SAN environment, please read Chapter 3, “Building a SAN environment” on page 51.

For information about how to set up the Storage Agent, read Section 5.3, “Installing and configuring the Storage Agent” on page 132. We describe all the necessary steps to prepare your environment for LAN-free data transfer.

5.5.1.1 Installing TDP for R/3

In this section we give you a brief overview of how to set up TDP for R/3 on a Windows NT system. For more detailed information, please read the documentation supplied with this product. You can also read the IBM Redbook, *R/3 Data Management Techniques Using Tivoli Storage Manager,* SG24-5743.

For our LAN-free tests, we used the following version of TDP for R/3:

```
Interface between SAPDBA Utilities and Tivoli Storage Manager
      - Version 3, Release 1, Level 0.2  for WinNT -
        Build: 110  compiled on Aug 30 2000
```

To install TDP for R/3 on your computer:

1. Run the distribution file **TDP_R3_V3.exe**. A screen as shown in Figure 65 will appear:

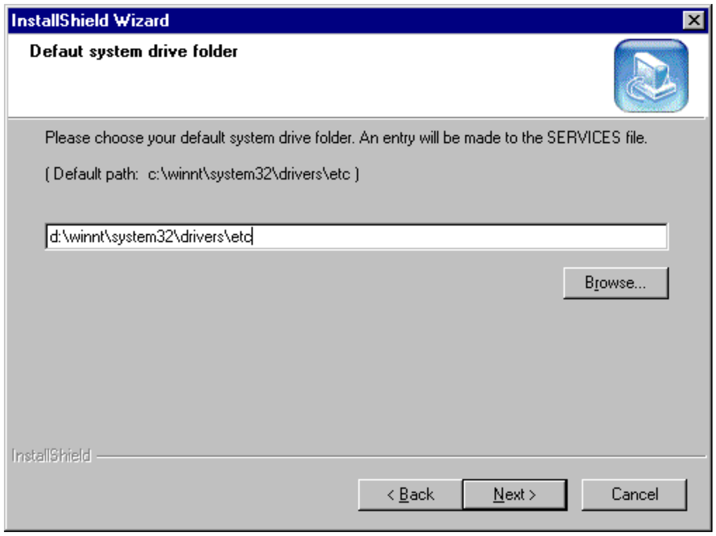

Figure 65. Choose path to the system folder

During the installation process, setup will ask you for the location of your SERVICES file on your NT box in order to add the following entry to this file:

```
backint    57321/tcp    #Tivoli Data Protection for R/3
```

There is a known bug in the installation routine; the installer will **not** modify the SERVICES file. You must manually add the line shown into this file. Otherwise, the service called PROLE will not start, and TDP for R/3 will not work. After you have modified the SERVICES file, reboot your machine.

2. Now specify the name of your Oracle instance, as shown in Figure 66.

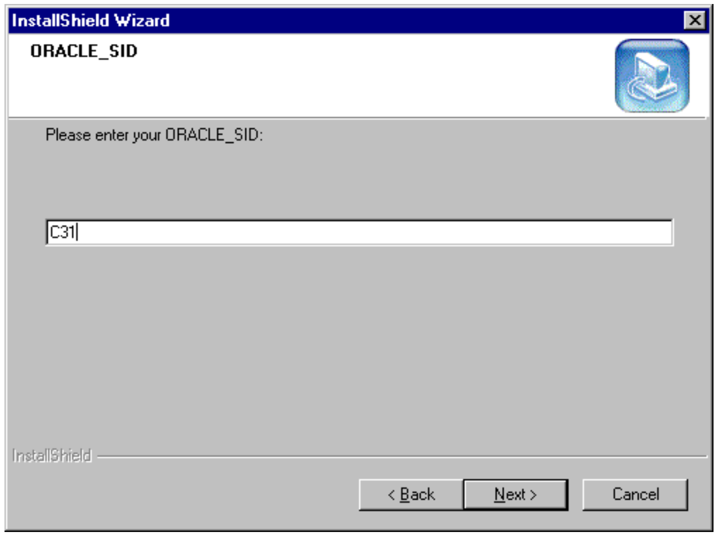

Figure 66. Specify ORACLE_SID

3. Enter your default system drive letter, as shown in Figure 67. The letter determines on which drive R/3 and TDP for R/3 will be installed.

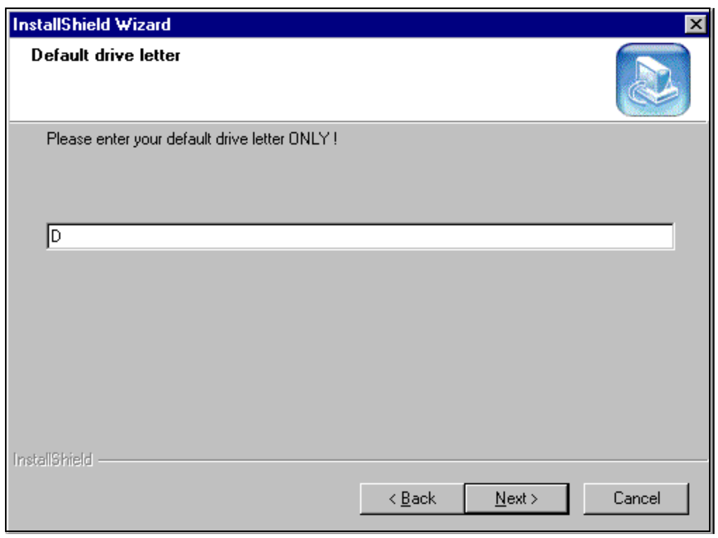

Figure 67. Enter the system default drive letter

4. As shown in Figure 68, now specify the version of your R/3 environment you have installed on your system. If you select the **Release 4.X** option, then the TDP's executable files will be by default installed to the <drive>:\usr\sap\<SID>\sys\exe\run folder.

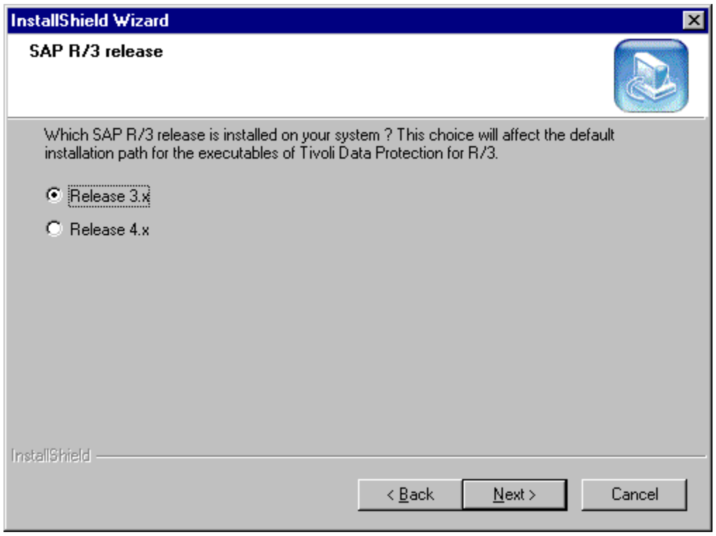

Figure 68. Specify the R/3 release

5. Select the standard type of installation (Figure 69)

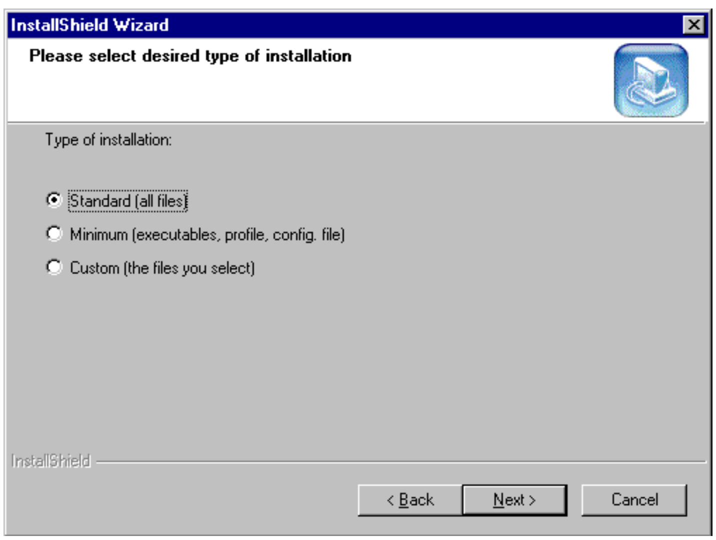

Figure 69. Selecting type of installation

6. Choose the destination folder, as shown in Figure 70, for all the TDP files (copy of the executables, sample scripts, documentation)

Figure 70. Choose destination folder

7. Choose the destination (see Figure 71) where you want to put the TDP executables and libraries needed for R/3 (`backint.exe`, `backfm.exe`, `prole.exe`, `backintv3_en_us.dll`). Check the path and correct it if needed.

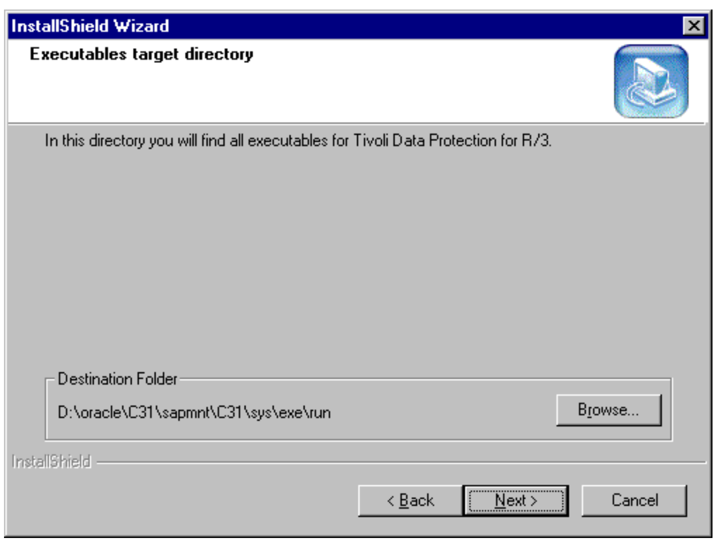

Figure 71. Choose folder for executables

8. Then, as shown in Figure 72, select the destination folder for the profile and configuration files `init<SID>.utl` and `init<SID>.bki`. It is very helpful to store these files in the same folder where your `init<SID>.sap` profile resides. In our case, is it the `d:\oracle\C31\dbs` folder.

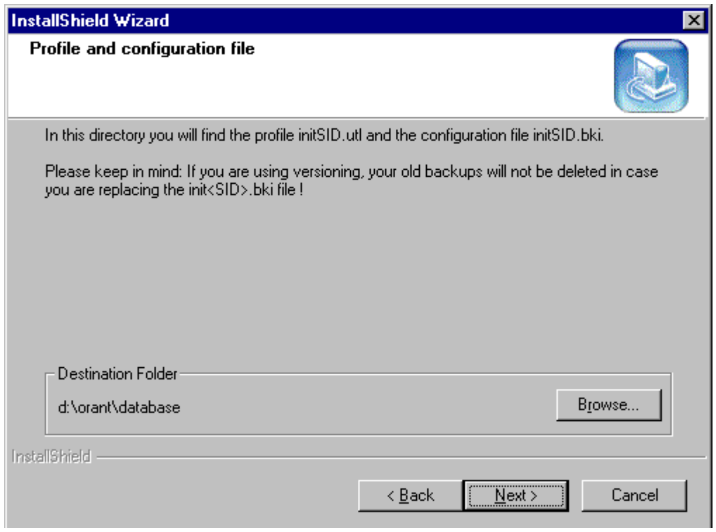

Figure 72. Choose destination for profile files

9. Check the destination where your `<servername>.opt` file(s) will reside and change it if needed (see Figure 73).

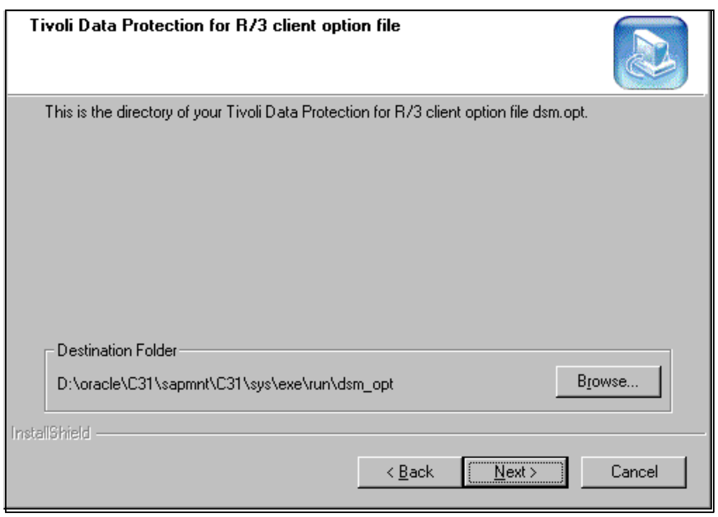

Figure 73. Option file destination folder

Then click the **NEXT** button, and the TDP files will be installed. At the end of the installation, the setup requires you to reboot the machine. Do *not* do this now. Before you reboot, modify the `\winnt\system32\drivers\etc\services` file as described in Step 1. Otherwise, the `PROLE` service will not start and you will not be able to set up TDP.

5.5.1.2 Configuring TDP for R/3

To configure TDP for R/3, you will need to go through the following steps:

- Set the environment variables for TDP.
- Modify the `init<SID>.sap` file.
- Modify the `init<SID>.utl` file.
- Create and modify the `<servername>.opt` file.

Setting environment variables

TDP for R/3 requires you to set the following variables:

- DSMI_CONFIG: This variable points to a `DSM.OPT` file.
- DSMI_DIR: This variable specifies the DSM API installation directory.
- DSMI_LOG: This variable points to a directory where the dsmierror.log file will be created.

Here is the example output of the `SET` command issued on our test system:

```
DSMI_CONFIG=D:\oracle\C31\sapmnt\C31\sys\exe\run\dsm_opt\dsm.opt
DSMI_DIR=D:\Program Files\Tivoli\TSM\api
DSMI_LOG=D:\temp\dsierror.log
```

Modifying the INIT<SID>.SAP file

In the `INIT<SID>.SAP` profile file, you need to declare that you will be using a “util file” through which your backup/restore commands (`BRBACKUP`, `BRARCHIVE`, `BRRESTORE`) can call the TDP to open a Tivoli Storage Manager session.

Add these two lines to your `INIT<SID>.SAP` file:

```
backup_dev_type = util_file
util_par_file = %ORACLE_HOME%\database\initC31.utl
```

Modifying the INIT<SID>.UTL file

In the `INIT<SID>.UTL` file you need to set BACKUPIDPREFIX parameter, which is according to your Oracle instance name. Then make sure that the CONFIG_FILE parameter points to the correct `INIT<SID>.BKI` file.

The other parameters are needed for Tivoli Storage Manager communication.

```
BACKUPIDPREFIX       C31___
CONFIG_FILE          D:\orant\DATABASE\initC31.bki
SERVER               brazil
SESSIONS             1
PASSWORDREQUIRED     no
ADSMNODE             C31
BRBACKUPMGTCLASS     STANDARD
BRARCHIVEMGTCLASS    STANDARD
```

Creating <SERVERNAME>.OPT file

Create a file named `<SERVERNAME>.OPT`, where the <SERVERNAME> is the same string you have set for the SERVER parameter (see the screen shown) in the `INIT<SID>.UTL` file. In our case, we have created the file named `BRAZIL.OPT`. See the following example content of this file.

Please note that the `DSM.OPT` file will for the purpose of the TDP remain blank. TDP only checks if the DSMI_CONFIG variable is set and that the file `DSM.OPT` exists.

```
COMMmethod         TCPIP
NODEname           C31
TCPPort            1500
TCPServeraddress   193.1.1.11
TCPBUFFSIZE        31
TCPWINDOWSIZE      32
ENABLELANFREE     YES
```

The ENABLELANFREE parameter enables TDP to communicate with the Storage Agent. This option is needed if you want to use SAN to back up your R/3 data.

5.5.2 Backing up and restoring R/3

In this section we will describe how to back up and restore your SAP system. Please keep in mind that, for our tests, we use only basic configuration and scenarios. If you are looking for more complex information about R/3 data management, we recommend that you read the IBM Redbook, *R/3 Data Management Techniques Using Tivoli Storage Manager*, SG24-5743.

5.5.2.1 Backing up R/3 using LAN-free data transfer

To back up our R/3 system, we used the following command:

```
brbackup -u system/manager -t online -c -m all
```

Wait until BRBACKUP calls the BACKINT program, which is the interface between Tivoli Storage Manager (in our case, the Storage Agent) and R/3. This will be indicated by the following display:

```
BR280I Time stamp 2000-09-27 13.03.24
BR229I Calling backup utility...
BR278I Command output of 'd:\oracle\C31\sapmnt\C31\SYS\exe\run\backint -u C31 -f
backup -i
d:\oracle\C31\sapbackup\.bddqvszh.lst -t file -p d:\orant\database\initC31.utl -c':

                         Tivoli Data Protection for R/3

         Interface between SAPDBA Utilities and Tivoli Storage Manager
               - Version 3, Release 1, Level 0.2  for WinNT -
                 Build: 110  compiled on Aug 30 2000
        (c) Copyright IBM Corporation, 1996, 2000, All Rights Reserved.

BKI0005I: Start of backint program at: 09/27/00 13:03:24 .
```

If you encounter an error similar to the one shown in the following screen, check if the PROLE service is started. If not, try to start it. In case you are unable to start the service, you probably have not modified your SERVICES file, as described in Section “To install TDP for R/3 on your computer:” on page 150.

```
BKI0005I: Start of backint program at: 09/27/00 13:09:58 .
BKI2001E: Socket error while connecting to PROLE on port backint: Unknown error.
```

If the BACKINT program is called and started successfully, you are now be able to see that the TDP session is started to the Storage Agent. You can verify this if you connect with the administrative command line client to the Storage Agent.

The next screen shows the output of the QUERY SESSION command issued in the Storage Agent. You can see that there are four TCP/IP sessions opened for the Storage Agent (in our example session numbers 49, 52, 56 and 58). These are used for accessing Tivoli Storage Manager server's database and library sharing.

Note there is also an TDP R3 session opened for client node C31 (session number 55). This is a named pipe session which has the TDP opened to the Storage Agent and through which sends the R/3 data. The circles indicate the client session opened for the Storage Agent and the actual amount of data transferred thought the SAN.

```
tsm: SA>q ses

  Sess  Comm.   Sess      Wait    Bytes    Bytes  Sess   Platform  Client
Number  Method  State     Time     Sent    Recvd  Type
------  ------  ------  ------  -------  -------  -----  --------  -------
    49  Tcp/Ip  Start      0 S    9.7 K   17.1 K  Serv-            SA
                                                   er
    51  Tcp/Ip  Run        0 S   57.3 K    1.5 K  Admin  WinNT     ADMIN
    52  Tcp/Ip  Start      0 S    9.4 K    9.9 K  Serv-            SA
                                                   er
    55  Named   RecvW      0 S      502  172.7 M  Node   TDP R3    C31
         Pipe                                             Intel
                                                         NT
    56  Tcp/Ip  Start      0 S  100.2 K   89.1 K  Serv-            SA
                                                   er
    58  Tcp/Ip  IdleW      0 S    9.6 K      166  Serv-            SA
                                                   er
```

If you need to backup/restore your data through the LAN, there are two methods you can use to do this:

1. Stop the Storage Agent service.

   or

2. Comment out the ENABLELANFREE option in the client option file.

If you then start a backup (restore), you can find out that there is no TDP session opened to the Storage Agent, but the session is opened through the LAN (TCP/IP) directly to the Tivoli Storage Manager Server.

In the next screen we issued the QUERY SESSION command on the Tivoli Storage Manager server. You can see there is a session opened for node C31 (session number 198), which uses TCP/IP transfer and 60 MB of data were already sent. In this case, LAN data transfer is being used.

```
tsm: BRAZIL>q sess

  Sess Comm.  Sess     Wait   Bytes   Bytes Sess  Platform Client Name
Number Method State    Time    Sent   Recvd Type
------ ------ ------ ------ ------- ------- ----- -------- --------------------
    26 ShMem  Run      0 S    7.9 K     685 Admin AIX      ADMIN
   185 Tcp/Ip IdleW  2.5 M   14.0 K  12.5 K Serv- Windows  SA
                                             er    NT
   188 Tcp/Ip IdleW    4 S  170.0 K 189.9 K Serv- Windows  SA
                                             er    NT
   190 Tcp/Ip IdleW    4 S      166  17.0 K Serv- Windows  SA
                                             er    NT
   198 Tcp/Ip RecvW    0 S    3.9 K  60.0 M Node  TDP R3   C31
                                                   Intel
                                                   NT
```

You can still access your data that was previously backed-up through the SAN via the LAN (TCP/IP), if needed. But be aware that data which you back up through the LAN will also be restored via the LAN.

5.5.2.2 Restoring R/3 using LAN-free data transfer

On our system, we only used the SAPDBA interface to restore the table PSAPUSER1D. We did this only to test the functionality of the LAN-free data transfer. The behavior of the Storage Agent, Tivoli Storage Manager server is basically the same for the backup and restore functions.

5.6 LAN-free client data transfer considerations

What happens if the SAN connection is now broken? In this case, the Storage Agent will be unable to communicate with the library and will stop sending data. The session will be in idle state until the SAN connection is available again. Then the data transfer will continue without any problems. If the SAN connection is **not** available longer than the value set for the COMMTIMEOUT parameter, the session will fail.

If the SAN connection is not available for the Storage Agent, and you start the backup via TDP, the Storage Agent will send a request for mounting a tape to the Tivoli Storage Manager server (in fact, to the library manager); the tape will be mounted; and the session will wait until the Storage Agent is able to communicate via SAN.

That means, if for any reason you want to back up your data through the LAN, you must stop the Storage Agent before you start the TDP client. If you do so, the TDP client will open a LAN connection to the Tivoli Storage Manager server, and the LAN transfer will be used. That is also the reason why there should be TCP/IP communication defined in the `DSM.OPT` file for TDP, as described in Section 5.4.2, “Installing and configuring TDP for MS Exchange” on page 141.

Keep in mind that the Storage Agent does **not** support disk storage devices for LAN-free transfer. This means that, for the destination in the copy group, if you choose a disk storage pool, the data will be backed up through the LAN.

In Figure 74 we choose to back up our data from the client machine to the disk storage pool on the Tivoli Storage Manager server. The client is configured for LAN-free data transfer. But because the Storage Agent does not support disk storage pools as a destination for LAN-free transfer, the client will send data across the LAN. Then, a server migration process occurs, and the data will be migrated from the disk storage pool volumes to the library through the SAN. When you restore the data, the Tivoli Storage Manager sever communicates with the library across the SAN, but the data will be transferred via the LAN.

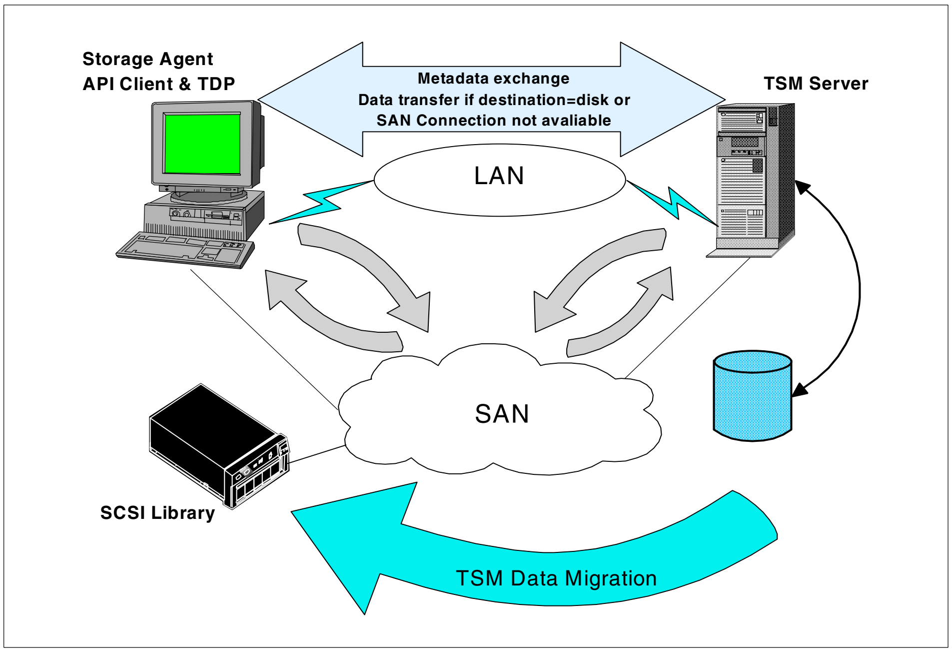

Figure 74. LAN/SAN backup/restore scenario

That is because the data has been originally backed up by the Storage Agent to the Tivoli Storage Manager server through the LAN. Such data cannot be restored through the SAN, even though the data has been actually sent from Tivoli Storage Manager server to the library via the SAN.

What happens if your LAN is not available before starting a backup/restore session on the client? The Storage Agent needs to access your Tivoli Storage Manager server's database, and this can only be done be through the LAN, so in this case, you will not be able to backup or restore any data.

And what happens if your LAN fails during the backup/restore? This depends on the exact moment when the LAN fails:

- Before the Tivoli Storage Manager server mounts a tape
- After the Tivoli Storage Manager server has mounted a tape

When your LAN fails *before* a tape is mounted on the Tivoli Storage Manager server, the backup/restore operation will fail with the error "Media mount not possible".

If the LAN fails *during* the backup/restore, its behavior is similar to what you would see when a regular client session loses its connection to the Tivoli Storage Manager server. The Client/Storage Agent will try to reopen the session; if it cannot, the backup/restore operation will fail.

For example, the R/3's utility, brbackup, will stop the current operation, and an error message appears stating that it cannot communicate with the Tivoli Storage Manager server. At this point, brbackup lets you choose whether to stop the brbackup or try to repeat the last failed action and continue. We waited until the LAN was able to operate again and let the brbackup continue. The brbackup resumed operation without any problems.

On the other hand, TDP for MS Exchange does not have a feature similar to the R/3's brbackup. Rather, TDP will not let you do anything, and just ends with an error, saying that it is unable to communicate with the server. You must start your backup again.

At this point, you may be also want to read about availability considerations in Chapter 8, "Availability considerations in a SAN" on page 235.

5.7 LAN-free file backup considerations

At the time of writing, you can NOT use the Backup/Archive (BA) client with the Storage Agent. The Storage Agent can be used only for API clients like Tivoli Data Protection products. So, currently, there are only limited possibilities for transferring your files with the BA client over the SAN.

One such method is to use library sharing. The disadvantage of such a solution is that you are required to have a Tivoli Storage Manager server installed on the client machine in order to act as a library client. In Chapter 4, “Tape library sharing” on page 75, we covered all the library sharing considerations.

Another way is to implement the SANergy File Sharing software, which gives you the possibility to backup/restore the data over the SAN locally on the MDC machine. Or you can backup the data to a locally installed Tivoli Storage Manager server on a SANergy client machine which mounts its volumes from the MDC machine.

You can read about SANergy concepts and usage in the Tivoli Storage Manager environment in Chapter 7, “Serverless backup using TSM and SANergy” on page 221. Another excellent source covering the theory of using SANergy in the Tivoli Storage Manager environment is the IBM Redbook, *Tivoli Storage Manager 3.7.3 & 4.1 Technical Guide,* SG24-6110.

Another possible solution would be to configure the TCP/IP communication over the SAN. This is supported, for example, by the SAN Fibre Channel 2109 switches and Qlogic’s 2200 series Host Bus Adapters. However, although this solution can help you offload the LAN traffic to the SAN fabric, keep in mind that this is not using the true concept of the LAN-free and server-free data transfer, which is based on direct access to storage devices.

Chapter 6. Migrating a TSM installation to a SAN environment

This chapter covers different scenarios and considerations to observe when migrating a classic Tivoli Storage Manager environment with local storage into a SAN environment. It emphasizes migrations involving server storage consolidation and Tivoli Data Protection for Application clients.

6.1 Differences in SAN and non-SAN TSM setups

The current implementations of Tivoli Storage Manager environments in a SAN environment, mainly consist of another way of attaching the SCSI tape libraries. Whereas classic installations use SCSI cable connections and native SCSI protocols to communicate with the library, the SAN installations instead use the SAN fabric to establish this communication. The SCSI library will be attached to the host using Fibre Channel with SCSI protocols, encapsulated in FCP packets. The result is that the main difference is located in the way the hardware components are connected to each other.

From a Tivoli Storage Manager point of view, the only link to the physical hardware is at the library and drive definition level. The storage volumes seen by the server and clients are logically linked to these definitions by storage pools and device class definitions. This means that the server sees these storage volumes independently of the hardware physically providing this storage.

These considerations allow us to conclude that there is no important difference between SAN and direct SCSI implementations of Tivoli Storage Manager. Even when library sharing is implemented, the only difference is that again, the physical definitions of drives and libraries will change.

6.2 Steps for migrating TSM into a SAN environment

Although no differences exist in the way storage volumes are seen by Tivoli Storage Manager, there are actions required when moving a current Tivoli Storage Manager implementation into a SAN environment. The main actions that will be required involve the following processes:

- Movement of volumes containing data, both logically and physically
- Server movement, meaning moving the Tivoli Storage Manager server definitions from one system to another
- Device redefinition

Besides these requirements, certain preparation tasks are needed, both at the hardware and software levels, and cleanup is necessary after the migration.

Three major migration scenarios will be described for three starting configurations (every scenario will consist of several sub-scenarios, depending on the required final configuration):

1. Migrating multiple Tivoli Storage Manager servers with locally attached tape storage to a shared SAN attached tape library.

   This will be done for library devices that both contain storage pools used directly by nodes, and storage pools receiving their data from other storage pools (migration). The main action in this case will be moving the storage pool volumes to a new storage pool. See Section 6.2.1.1, “Migration-receiving storage pools” on page 168, and Section 6.2.1.2, “Directly node-attached sequential storage pools” on page 173.

2. Migrating TDP for application clients from LAN environments to SAN environments.

   In this case, no major changes are required, but additional software installation (Storage Agent) and configuration will be required. See Section 6.2.2, “LAN attached application servers” on page 175.

3. Migrating TDP for application clients together with a local Tivoli Storage Manager server and tape library, to a central Tivoli Storage Manager server with a SAN attached tape library.

   In this case, the main tasks will involve server movement and library movement. Also, installation of the Storage Agent will be needed. Distinctions will be made between actual physical movement of the tape library, with equal or different server platforms between the original and the new situation, and movement of the data to a new library. See Section 6.2.3.1, “Moving the library to an equal operating system” on page 181, Section 6.2.3.2, “Moving the library between different server platforms” on page 184, and Section 6.2.3.3, “Using a new library” on page 188.

6.2.1 Consolidation of multiple TSM server attached libraries

Our first migration scenario will describe the consolidation of multiple locally attached SCSI tape libraries to a single shared library. The main reason for such a consolidation is to create a more efficient tape subsystem for the Tivoli Storage Manager server (see Section 2.2.3, “Efficiency” on page 43), and reduce the overall cost of the implementation. The initial situation is shown in Figure 75. This implementation consists of several Tivoli Storage Manager server entities, each with its own dedicated storage (both disk and tape).

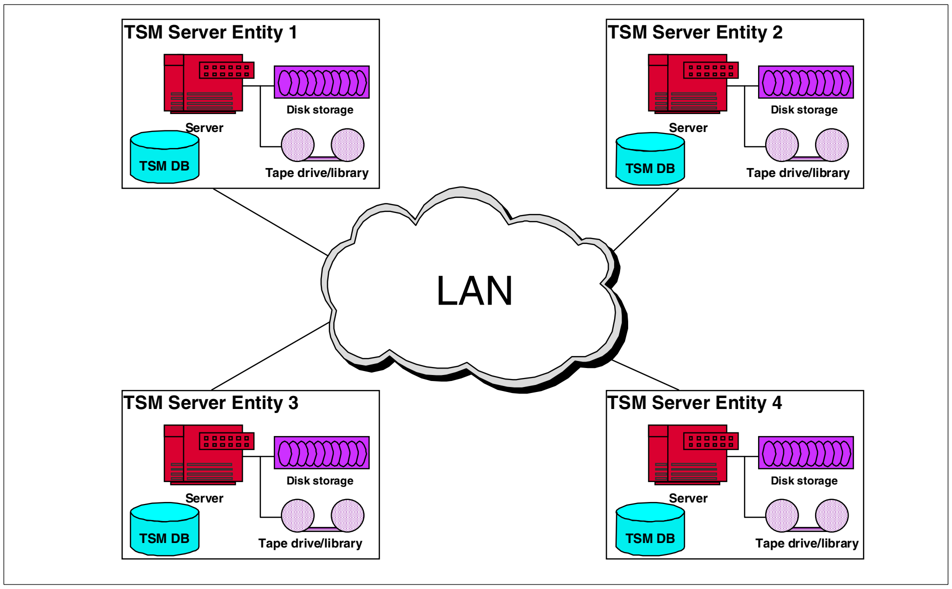

Figure 75. Initial server configuration

Figure 76 shows the goal environment for the migration. The same Tivoli Storage Manager server entities exist with their own database and disk storage, but with shared, SAN connected tape storage. This shared tape library is under control of one Tivoli Storage Manager server, acting as the library manager. Server-to-server communication (over the LAN) is used between this library manager and the other servers, or library clients, in order to obtain tape volumes. Data moves directly over the SAN from each server to the tape library. For more details on how to set up library sharing, please refer to Chapter 4, “Tape library sharing” on page 75.

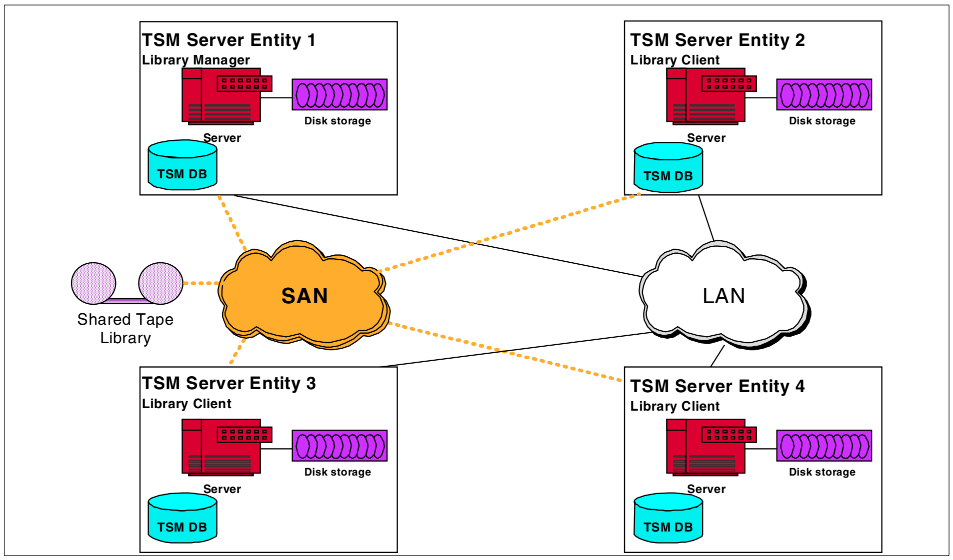

Figure 76. Shared library configuration

The major steps will include moving the storage pools and associated volumes belonging to the local library to the new shared library. The assumption is made that this shared library is available at the time of migration.

In case you intend to reuse one of the local libraries as shared library, it is important to move this one first. As indicated in Section 6.1, “Differences in SAN and non-SAN TSM setups” on page 163, this involves installing and defining the necessary hardware, as well as updating the library and drive definitions on the server. Also, you need to define the server, which is currently using the library, as the library manager. Figure 77 illustrates this operation. If you will be using a new library, the only step required is to connect this library to the server which will act as the library manager, and define this server as library manager. To move the data, follow these steps.

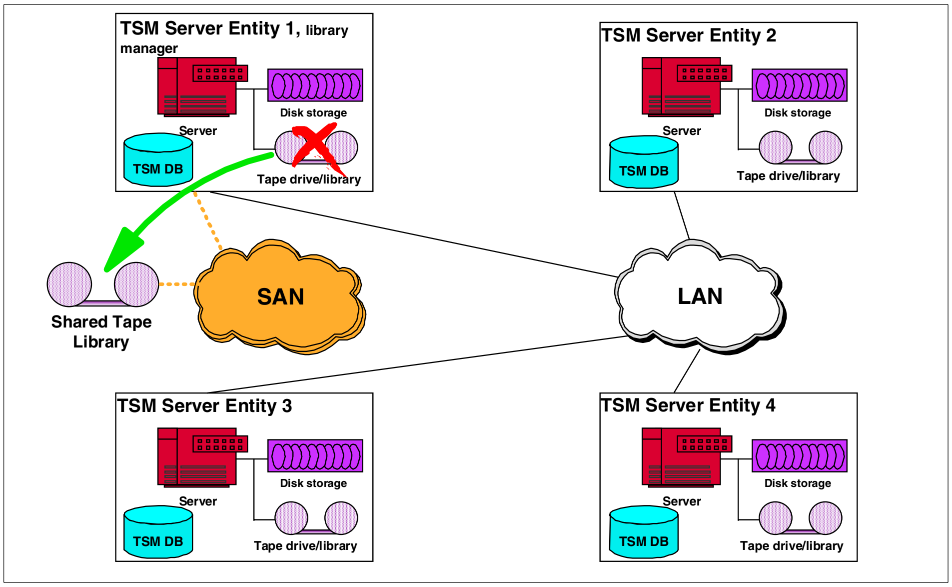

Figure 77. Moving the future shared library

As stated before, the main action in this situation will be to move data. However, there is also a need to update server definitions so that future operations will send their data to the new primary storage pools, instead of the original ones. Therefore, we will distinguish the following situation:

- Devices with tape storage pools that receive their data from another storage pool using migration processes
- Devices with tape storage pools that receive their data directly from a node, as defined in the used copy group for that node

Before starting the migration, the first requirement is to find out which storage pools will be eligible for migration. The pools that can be migrated are those associated with the device class pointing to the library that you will move. The second requirement is that these pools are primary sequential storage pools. Copy pools are not candidates for this migration. This means that you will have to delete the original copy pools, and recreate a copy pool using the new shared library device class. There is one exception to this situation, and that is the case in which you will move to a library which uses media compatible with the one you were using previously. In that case, it is enough that you update the device class used by the copy storage pool, so that it points to the new library, or update the physical library definition, so that is uses the new shared devices.

The reason for this need to recreate storage pools follows from what they are: a copy of the primary storage pools. This means that, in case you move them, the results and the time required will be exactly the same as the time needed when creating a new backup.

These considerations bring us to a last very important point. The movement of data using the steps described in Section 6.2.1.1, “Migration-receiving storage pools” on page 168, and Section 6.2.1.2, “Directly node-attached sequential storage pools” on page 173, are only required when the new library format is incompatible with the one that is currently used. If both libraries use compatible media, the only operation required would be to update the device class or library definitions, so that they point to the new library. The data movement simply consists of physically moving all volumes to the new library, check them in on the library manager.

6.2.1.1 Migration-receiving storage pools

This situation will explain the necessary steps to migrate data from a primary sequential storage pool (destination storage pool in Figure 78) to a new storage pool, using a SAN attached shared library. The storage pool receives its data using migration processes. This means that nodes do not send their data directly in this storage pool (except for situations in which the MAXSIZE parameter is in effect). The initial situation is shown in Figure 78.

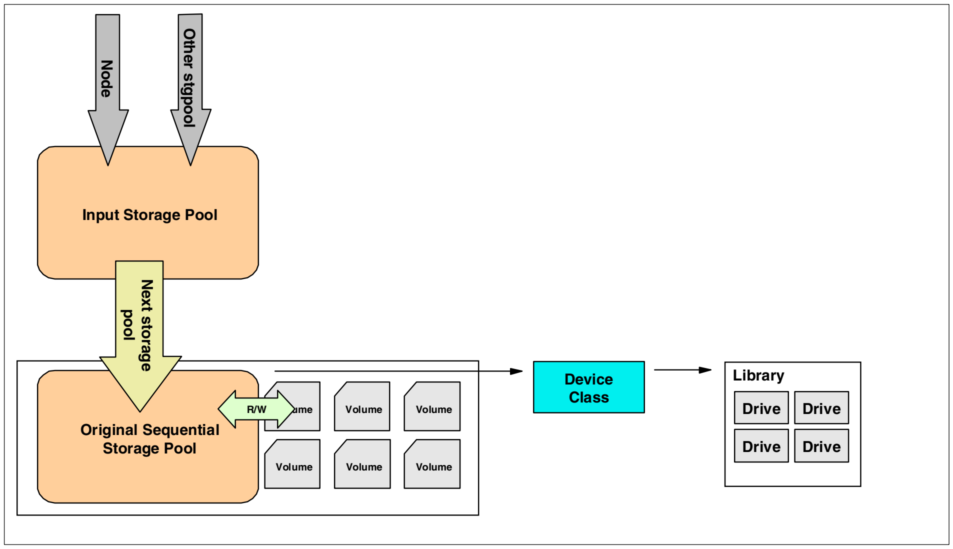

Figure 78. Initial situation

1. The first phase in the migration will prepare the Tivoli Storage Manager server for the migration. Its main purpose is to disable all data movement between storage pools, and to create an image of the server just before the migration (see Figure 79).

   a. Create a database backup. You can use a full or snapshot type database backup.

   b. Disable the server for client operations (1). Also, disable all administrative schedules that might have an influence on the storage pools (for example, database backup, migration, reclamation, and storage pool backup processes).

   c. For each of the primary tape storage pools that require migration, perform the following steps (numbers refer to the diagram):

      - Update the input storage pool so that it will not migrate. This can be done by setting the HIGHMIG value of the disk storage pool to 100. Also, update this value on all storage pools that might migrate to the input storage pool (2).

      - Update the original tape storage pool, setting all volumes to read only (3), and the MAXSCRATCH parameter to zero. Also set the REUSEDELAY parameter to a value that will allow you to reuse the database backup created in step 1a on page 169. Although both are not directly related, setting the REUSEDELAY parameter will allow volumes to be available when the database backup is restored. This also means that when volumes have a status of empty (volumes will become empty by moving the data to the new storage pool), they will not return to the scratch volume pool (if applicable). The number of days you want to keep your volumes and database backup should be equal to the estimated validation time for your new environment.

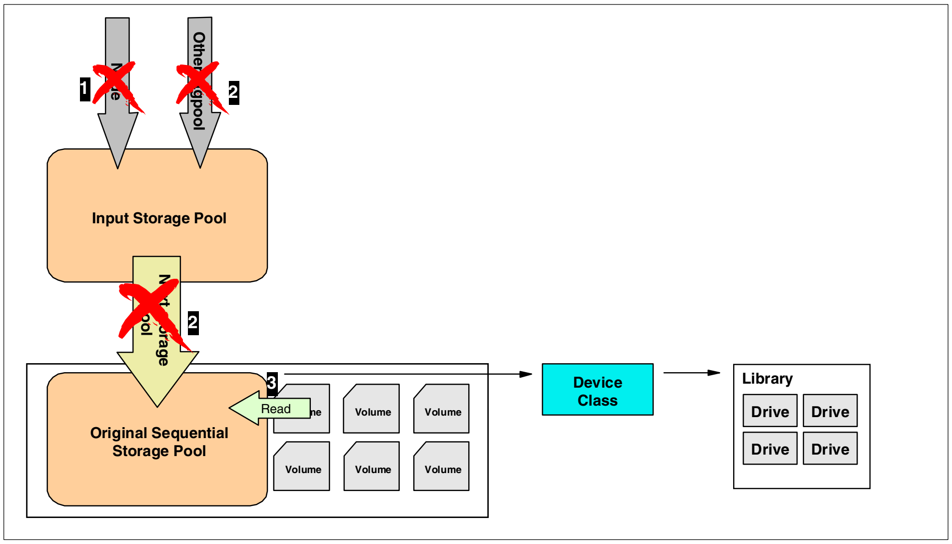

Figure 79. Preparing server storage pools before migration

2. Next, you need to install the required hardware and system software to attach the Tivoli Storage Manager server to the SAN.
3. After preparing your current storage pools and setting up the hardware, the next step will be to define the new devices, device classes, and storage pools (see Figure 80).
   a. Define the shared library and drives (4). If this migration is the first one, you should also define this server as library manager.
   b. Define a new device class, using the tape library defined in (5).
   c. Define one or more storage pools (6), using the device class defined in step 3b. You could use the same parameters as the ones used for the original sequential storage pools. If you will use private volumes for the storage pools, define them also.
   d. If this is the library manager, check in new scratch tape volumes into the library.

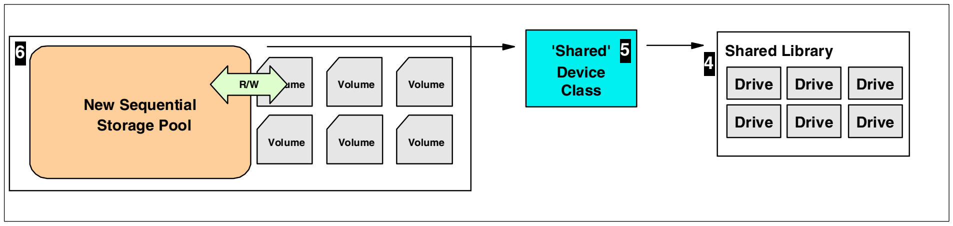

Figure 80. Defining new devices

4. After setting up the new storage definitions, it is now time to move the data from the original storage pools to the storage pools defined in the shared library. The steps in this section should be repeated for all storage pools that are moved (see Figure 81).

   a. Update the original storage pool, setting the newly defined one as its next pool (7), using the NEXTSTORAGEPOOL value.

   b. Set the HIGHMIG and LOWMIG values of the original storage pool to 0. This action will trigger a migration process, moving all data from the original pool to the pool in the shared library (8).

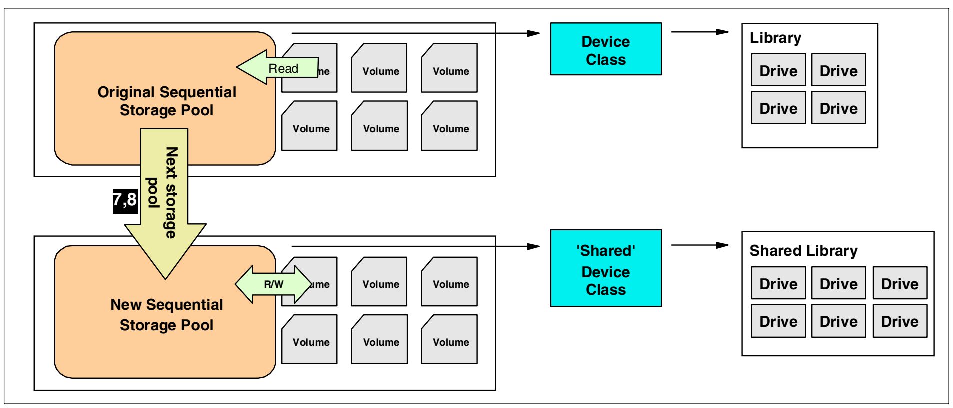

Figure 81. Moving data

   c. After the migration process completes, update the original input storage pool so that it uses the newly defined storage pool as next storage pool (see Figure 82). Also reset the HIGHMIG and LOWMIG values to their original values (9) (they were set to 100 in step • on page 169).

5. Re-enable the server for client sessions (10), unlock previously locked administrators, and reactivate the administrative schedules.

6. After completing these steps, your server is ready for operations. The only thing that is still necessary is to clean up all server objects that are no longer used. These include the original library and drives, the device class, the storage pools, and the volumes (11).
   a. Delete volumes belonging to original storage pools.
   b. Delete original sequential storage pools (including copy storage pools).
   c. Delete original sequential device classes.
   d. Delete original tape drives and libraries.
   e. Modify all references (like administrative schedules) to the original storage pool, and have them pointed to the new one. Also redefine a new copy storage pool, and update the storage pool backup schedule.

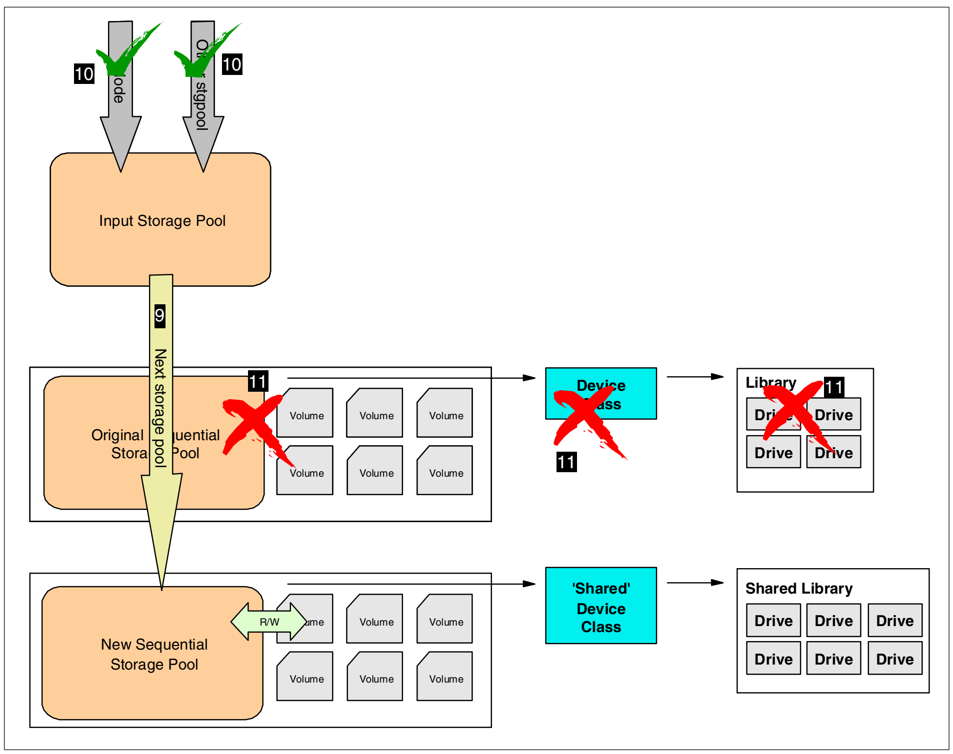

Figure 82. Finalizing definitions and cleanup

6.2.1.2 Directly node-attached sequential storage pools

This situation will explain the necessary steps to migrate data from a primary sequential storage pool to a new primary sequential storage pool, using a SAN attached shared library. The storage pool receives its data directly from a node, as defined in the copy group of the management class the node is using. Figure 83 shows the initial situation.

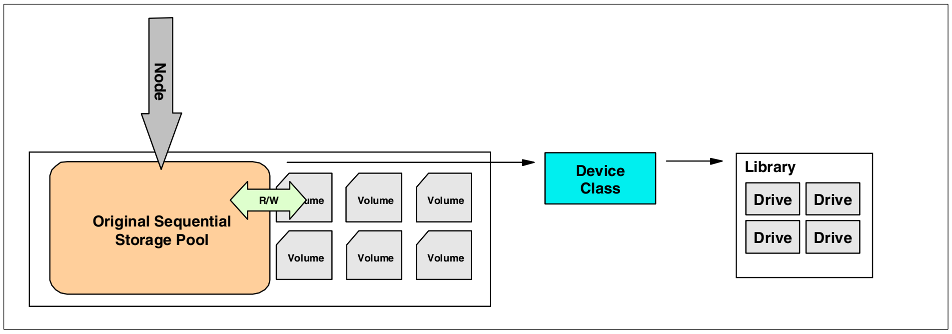

Figure 83. Initial situation

1. The server preparation will be the same as in the previous scenario (see step 1 on page 169). The only difference is that since the sequential storage pool will not receive data from another primary storage pool using the migration, you do not have to update input storage pools within the storage hierarchy.

   As shown in Figure 84, the actions required include disabling the server for client sessions (1) and setting the volumes to read only (2).

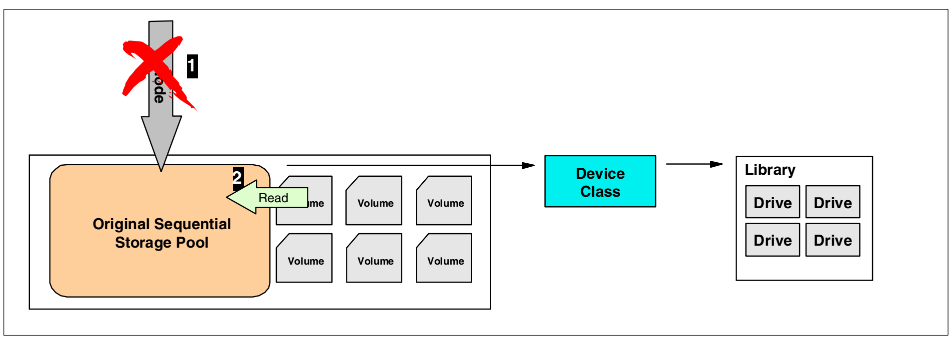

Figure 84. Preparing the server

2. Steps 2 and 3 are the same as in the previous situation (see steps 2 and 3 on page 170).
3. After defining the new shared library device, device class and storage pool, we must move the data. As in the previous situation, we will use a migration process to move data from the original primary storage pool to the newly defined primary storage pool on the shared library. To trigger the migration, use the following steps.
   a. Update the original sequential storage pool by setting the newly defined one as its next pool within the storage hierarchy.
   b. Set the HIGHMIG and LOWMIG values of the storage pool to 0. This action will trigger a migration process, moving all data from the original pool to the pool in the shared library.

Figure 85 shows the next steps in the process. Since the way new data is entering the storage pool differs from the previous situation, we must change other definitions.

4. Nodes or certain node file types are associated with a management class. This management class uses copy groups, one for backup objects and one for archive objects. Within the copy group, the destination storage pool for backup and archive operations needs to be changed from the original storage pool (3) to the new storage pool (4). Also, if HSM is implemented, the MIGDEST value of the management class needs to be updated.
5. Re-enable the server for client sessions (5).

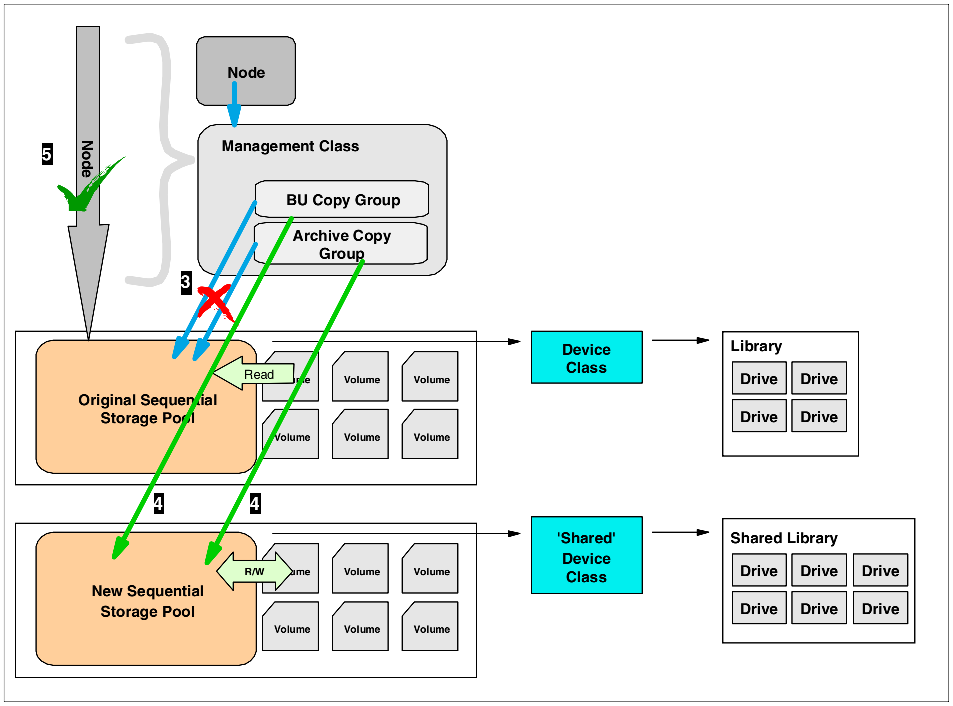

Figure 85. Updating copy groups

6. As in the previous situation, the migration is finalized by deleting all objects associated with the initial library, drives, device classes, storage pools, and storage pool volumes. Also, modify all references (such as administrative schedules) to the original storage pool, and have them point to the new one.

6.2.2 LAN attached application servers

A second type of installation that could be considered for SAN integration uses Tivoli Data Protection clients, backing up their data to a LAN Tivoli Storage Manager server. These types of application servers usually hold large amounts of data. Therefore, they tend to have performance problems due to LAN bandwidth limitations. A solution is to install the *Tivoli Storage Manager Storage Agent* (see Section 5.1.3, “The Storage Agent” on page 127), which will allow direct backups over the SAN network.

Currently, this solution exists for MS Exchange and SAP R/3 application servers running on Windows NT. One limitation that exists when using the Storage Agent, is that data needs to be backed up to tape volumes. Presently, no disk storage pools are supported. For details on the Storage Agent, please refer to Chapter 5, “LAN-free client data transfer” on page 125.

The starting point of this migration scenario is shown in Figure 86. Using TDP for Applications and the normal Tivoli Storage Manager Backup/Archive client, data is moved over the LAN to the Tivoli Storage Manager server. The Tivoli Storage Manager server has a locally attached SCSI tape library, or a library which is already attached through the SAN. The data that is backed up from the client is split into two parts: application data, which is under control of the application running on the application server, and other data, which typically consists of system files and user data. The application data will be backed up using the TDP component; the files or other data will be backed up using the Tivoli Storage Manager B/A client.

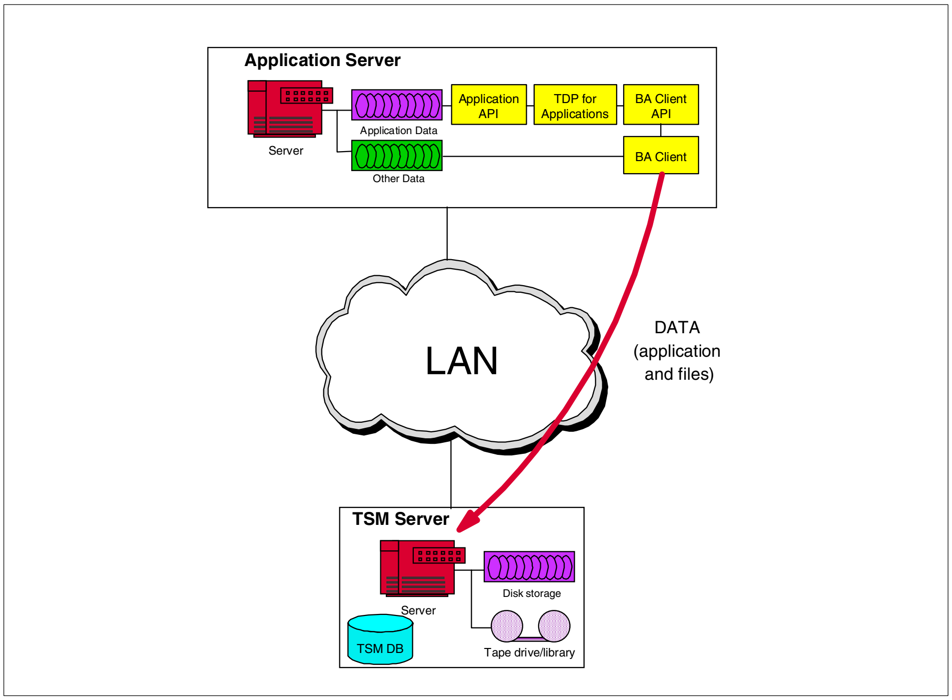

Figure 86. LAN attached application server

Figure 87 shows the configuration which is the goal of this migration. Application data is passed directly over the SAN to the SAN attached tape library. This is done using the Storage Agent, which communicates with the Tivoli Storage Manager server, sending and receiving metadata. The normal files are still transferred over the LAN network.

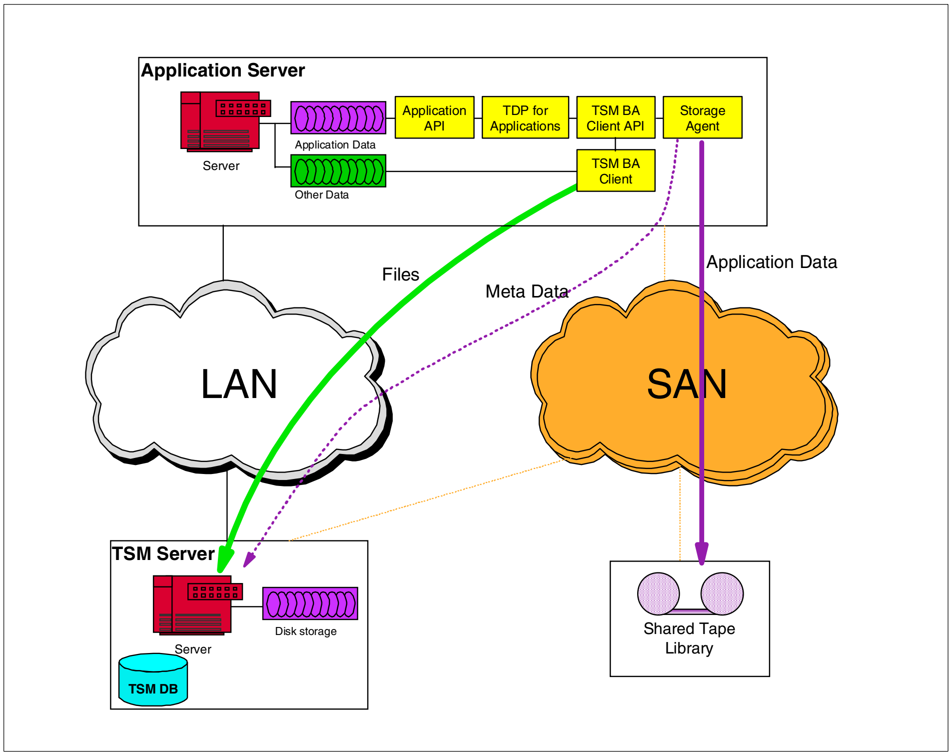

Figure 87. Application server using the Storage Agent

The migration scenario can be split into the following tasks:

- Library and tape device considerations
- Moving data
- Application server preparation
- Configuration of TDP
- Storage Agent installation and configuration

The following explains these tasks in more detail:

1. As shown in Figure 87, the tape library should be SAN attached to a Tivoli Storage Manager server, acting as library manager. The TDP client should send its data directly to a storage pool using this shared library. This setup is the same as the one explained in Section 6.2.1, “Consolidation of multiple TSM server attached libraries” on page 165. Attention should be paid to one thing: the Tivoli Storage Manager Backup/Archive client running on the application server does not use the LAN-free setup. It still sends its data over the network. Therefore, it is imperative that you have a least two management classes defined, one for the TDP client, sending the application data, and one for the Tivoli Storage Manager B/A client sending the other data. The setup of the management class and copy groups for the TDP client is described in Section 6.2.1.2, “Directly node-attached sequential storage pools” on page 173. The setup for the normal Tivoli Storage Manager B/A client, which will normally use a disk pool as initial target for its data, where it will be migrated to the tape pool, is described in Section 6.2.1.1, “Migration-receiving storage pools” on page 168.
2. The next step is to move the data residing on the TDP client’s original disk storage pools to the new tape pool. Since the configuration that is going to be created only supports tape storage pools, we must be certain that all TDP client data resides on this tape storage pool. If you were using a disk pool in the original setup, you must now move the data from that disk pool to the tape storage pool. This can be done using the `move data` command (3). You could also use a construct that sets the new tape storage pool as next storage pool for the disk pool, and migrate the data. However, since disk storage pools normally consist of a small number of volumes, moving the data on command might be faster. Figure 88 shows the results of steps 1 and 2.

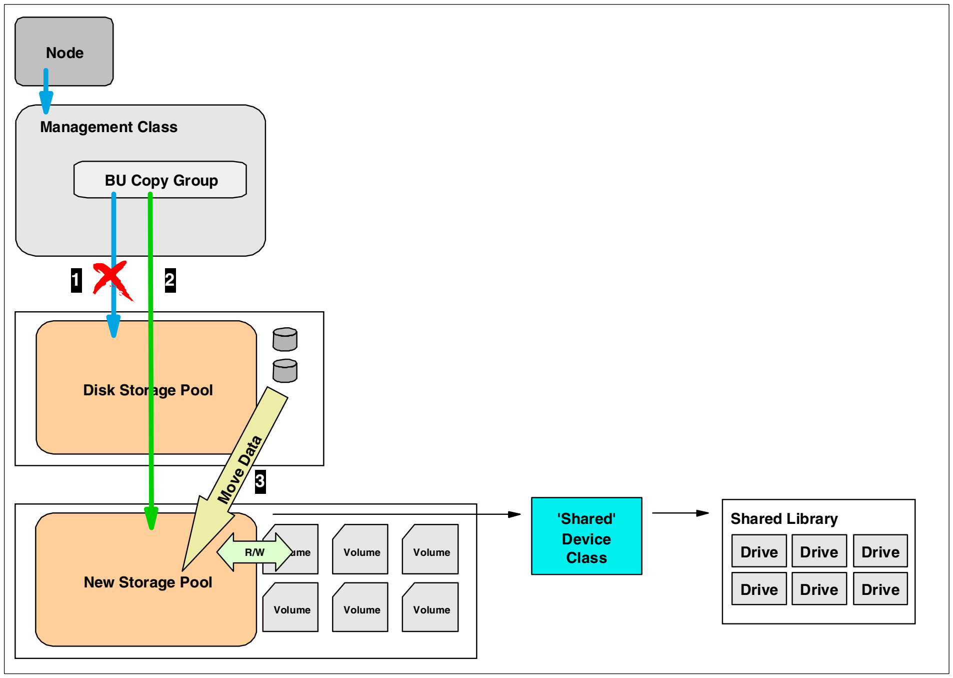

Figure 88. Redefining and moving data for TDP clients

3. After moving the data, the application server must be modified so that it can use the SAN fabric. You should install and configure all required hardware components.
4. Next, the TDP client and Storage Agent software must be configured and installed on the client. This involves the following steps:
   a. Install the Storage Agent software on the application server.
   b. Update the TDP client option file, adding the ENABLELANFREE YES statement.
   c. Modify the Storage Agent option file (dsmsta.opt), adding the name of the device configuration file that will be used.
   d. Issue the `DSMSTA SETSTORAGESERVER` command, which will generate the required server definition entries in the specified device configuration file.
   e. On the Tivoli Storage Manager server, define the Storage Agent as server, using the `DEFINE SERVER` command. Your Tivoli Storage Manager server should act as library manager.

f. Start the `ADSMSCSI.SYS` driver on the application server. This driver will be used by the Storage Agent to connect to the library.

g. On the Tivoli Storage Manager server, define drive mappings for the drives that will be used by the Storage Agent.

This completes the setup for this type of migration. For more detailed information about the installation and configuration of the Storage Agent, refer to Section 5.2, “Setting up LAN-free client data transfer” on page 128.

6.2.3 Locally attached tape library

A final configuration that is eligible for migration into the SAN is the application server, with locally installed Tivoli Storage Manager server and media. As stated in Section 6.2.2, “LAN attached application servers” on page 175, application servers tend to generate large amounts of data, which might not be able to pass through the normal network bandwidth. A solution many users have adopted is to directly attach a tape device or tape library to the application server, and run a local Tivoli Storage Manager and TDP client (see Figure 89).

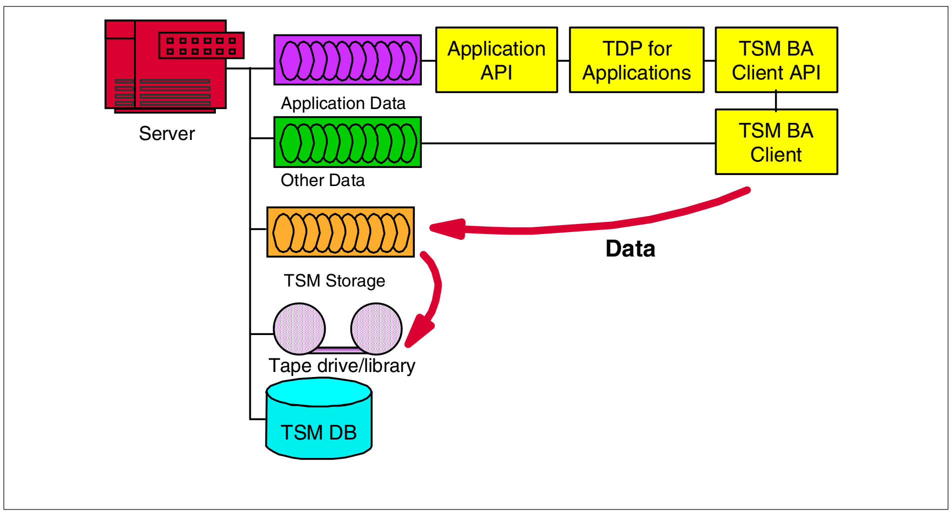

Figure 89. Application server and TSM server on one machine

Although this *stand-alone backup solution* solves the network problem, there are numerous disadvantages; the most important one is probably the use of system resources. Running two application servers, a database, and a storage manager on one system will increase the need for CPU, memory and disk space. A second disadvantage is that the machine holds all information required for database restores or recoveries. However, the backed up information is also under control of a database, the Tivoli Storage Manager server database, which resides on the same machine. This means that in case of system failure, the restores become much more complex. First, the entire Tivoli Storage Manager environment needs to be restored, and only after that is finished, can the application data be restored.

Of all the proposed migration scenarios, this one is the most complex, and will involve the most work. The reason for this is that the previous two installations did not require any Tivoli Storage Manager server movement. The tape volumes used before and after the migration were always under control of the same Tivoli Storage Manager server. In this case, the initial Tivoli Storage Manager server, and all of its storage, needs to be moved to a new machine. The resulting layout is the same as in the previous situation (see Section 6.2.2, “LAN attached application servers” on page 175), as shown in Figure 87 on page 177.

When considering this move, two scenarios can be pursued:

- Logically moving the current library and definitions to a new Tivoli Storage Manager server, attaching it physically to the SAN. The steps involved also depend on the chosen server platform for the central Tivoli Storage Manager server. The reason for this is that the logical move of the library to another Tivoli Storage Manager server means that we need to move the Tivoli Storage Manager server database. If the initial platform is the same as the new one, we can use the database backup and restore mechanisms. Otherwise, an export and import process will be required.
- Moving all data and definitions to an existing Tivoli Storage Manager server, with LAN attached tape storage. This means that the original library will not be reused.

Based on the options taken, the migration steps required will differ. This is explained in the following sections.

6.2.3.1 Moving the library to an equal operating system

In this case, we will move the library from a Windows NT system or Windows 2000 system to the same platform on the new Tivoli Storage Manager server. The following steps explain the required actions.

1. Create a backup of the original Tivoli Storage Manager system.
   a. On the original system, move all the data that resides in the disk storage pools to the tape pools (see Figure 90, Step 1).
   b. Create a Tivoli Storage Manager server database backup on the original system. The type of backup can be a full backup or snapshot (2). Also copy the volume history file, as well as the device configuration file (3). Make sure that the backup that you make is available after the move. The easiest solution is to use a volume that is currently in the library (or more than one volume).
   c. Finally, copy the server option file and the required license files.
   d. Stop the server.

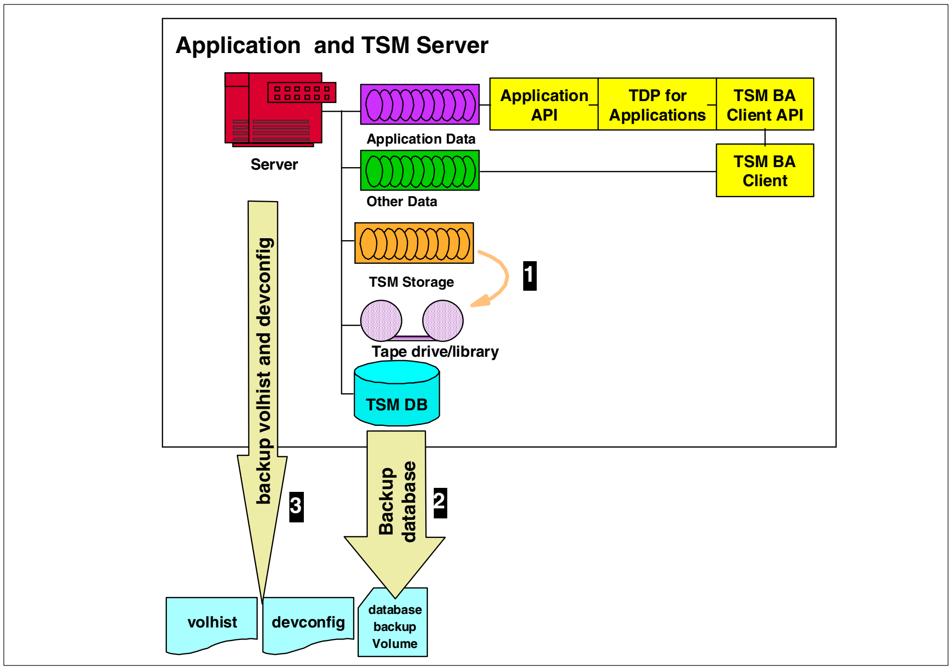

Figure 90. Backing up the initial configuration

2. Next, you should install the necessary hardware and software on the new machine, which is going to be used as Tivoli Storage Manager server.
   a. Attach the library through the SAN infrastructure to the new Tivoli Storage Manager server that you are about to configure.
   b. Install the Tivoli Storage Manager server.

3. Now, it is time to restore the configuration as it was on the original machine. On the newly installed Tivoli Storage Manager server, perform the following tasks:
   a. Start the `ADSMSCSI.SYS` device driver and check the device names for the library and drives of the SAN attached library.
   b. Update the device configuration file from the original machine, replacing the old device names with the new ones.
   c. Copy the modified device configuration file and volume history file to the server instance directory that you will use. Also recopy the server option file and the license files.
   d. Run a database restore, using the point in time mechanism. This backup will read the device configuration file and the volume history file to access volumes in the library. As a result of this operation, your database will be exactly the same as it was when you created the backup.
4. When you now start the server, you still have to modify the disk storage pool volumes. In the restored database, they pointed to volumes as they were installed on the original system. There is a good chance that the name, size, or location of these volumes has changed. Also, the destination copy group for the TDP client must be changed from the disk storage pool to the tape storage pool.
   a. Create new disk volumes for the disk storage pool. When establishing the required size, please keep in mind that these volumes will no longer be used by the TDP client. This is due to the fact that the TDP client will send its data immediately to the tape library, using the Storage Agent.
   b. Define these volumes to the disk storage pool, deleting the old volume names.
   c. Modify the copy group in the management class used by the TDP client, so that it points directly to the shared tape storage pool.
5. Finally, you should modify the configuration of the original application server.
   a. Remove the original Tivoli Storage Manager server. Since the Storage Agent cannot be used on a system where the Tivoli Storage Manager server is installed, this step is required.
   b. Update the Tivoli Storage Manager client option file, so that it points to the new server. This connection will be used for normal Backup/Archive data transfers.

c. For the TDP part, the actions required are the modification of the option file, installation and configuration of the Storage Agent and defining the Storage Agent to the server. To do this, follow step 4 on page 179.

6.2.3.2 Moving the library between different server platforms

A second possibility, when moving the library, is to select a Tivoli Storage Manager server that runs on a different platform than the original Tivoli Storage Manager server, running on NT or Windows 2000. The difference between the movement among servers running the same operating system (see Section 6.2.3.1, “Moving the library to an equal operating system” on page 181), is that you cannot simply back up and restore the database. The reason for this is that the database contains platform specific information. The supported way to move a server, including the definitions and the storage volumes, is to perform aa Tiv Tivoli Storage Manager server export and import operation.

The following steps need to be performed:

1. The first step will be the preparation of the Tivoli Storage Manager server running on the application server. The tasks involved are shown in Figure 91.
   a. Start by backing up the original Tivoli Storage Manager database (1). This will allow you to fallback to the state you were in before starting the migration. A full or snapshot database backup can be used for this purpose.
   b. Disable all client sessions and administrators, and inactivate all administrative schedules.
   c. Move all the backup data belonging to the TDP client and residing in a disk storage pool to the tape storage pool (2).
   d. After doing this, update the management class of the TDP client, so that it points to the tape storage pool (3). You can also already delete the disk storage pool used by the TDP client, and its associated volumes (4). The reason for doing this is that, as stated before, the Storage Agent will not use a disk storage pool. To simplify the import process afterwards, it is a good practice to set up this situation as it will be in the new environment.
   e. Also note the definitions of all devices (library and drives), device classes, storage pools, and disk storage pool volumes. This might be useful when they will need to be redefined on the new server.

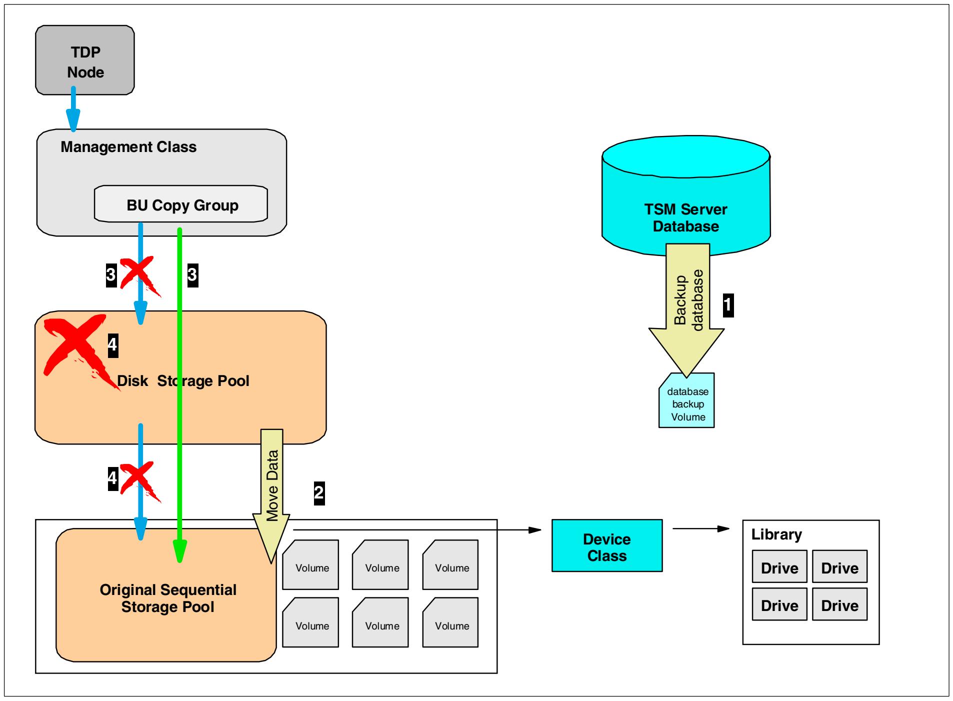

Figure 91. Preparing the server for export

2. After the preparation is done, you can export the server. Exporting information from a Tivoli Storage Manager server means that you will create sequential access volumes, containing all server definitions and data. The sequential volumes used should be accessible afterwards by the new server. If possible, the easiest way is to export them to the library that you are going to move.

   However, keep in mind that since you are exporting all data, you will need the same number of storage space for your export volumes as you will need for your current storage pools, and that you need at least two tape drives in the library. Also, you can use several levels of export, meaning that you can export the entire server, only the client definitions and data, the administrative definitions, or the policy definitions. In this case, the easiest way would be to export the entire server in one operation. This requires the following processes:

   a. Use the `EXPORT` command with the `PREVIEW` parameter to determine which objects and how much data will be moved. The server sends the following types of messages to the activity log for each operation:

- Reports on the types of objects
- Number of objects
- Number of bytes that would be copied to sequential media volumes

Use this information to determine how many sequential media volumes you will need.

b. If you would like to use specific labeled volumes to contain your export date, label them now. If this is not required, you can use the scratch pool to automatically get tapes.

c. Perform the actual export process, including the data. If you use the `EXPORT SERVER` command, the following objects and data will be exported:

- Administrator definitions
- Client node definitions
- Policy domain, policy set, management class, and copy group definitions
- Schedule definitions and client node associations
- File space definitions
- File space authorization rules
- Backed up, archived, and space-managed files

For the backup data, you have the possibility to only export the active versions of the files.

3. Next, you should install the necessary hardware and software on the new machine, that is going to be used as Tivoli Storage Manager server.

a. Attach the library through the SAN infrastructure to the new Tivoli Storage Manager server that you are about to configure.

b. Install the Tivoli Storage Manager server. This step includes defining disk space for the database and recovery log, as well as setting the server options and registering the licenses.

c. Define server storage for the target server.

Because each server operating system handles devices differently, Tivoli Storage Manager does not export server storage definitions. Therefore, you must define initial server storage for the target server. This means that you must define the shared library, drives, device classes and storage pools. Also, if a disk storage pool is to be used, you should define and format the disk volumes that will be used.

4. After the base configuration of the Tivoli Storage Manager server is completed, you can start the import process using the export media created earlier on. Be sure that you have access to this media.

a. By starting with running the import process with the preview option, you will get information in the server activity log concerning what and how much data will be imported.

b. Next, import server control information, which includes administrator definitions, client node definitions, policy domain, policy set, management class, copy group definitions, schedule definitions and client node associations.

> **Note**
>
> By default, the import process does not redefine objects that are already present on the server. If you use default names for some objects, you might want to consider doing the import process with the replacement parameter set to yes. Also, when importing several servers on one machine, you could encounter situations where you use the same name for certain objects on all servers, but with different attributes. In that case, you should rename these objects after the import, so that subsequent imports will not fail or redefine the original ones.

c. After the import process, you can check if all storage definitions comply with the destinations specified in the copy groups. To do this, you can use the following techniques:

   - If the policy definitions you imported included an ACTIVE policy set, that policy set is validated and activated on the target server. Error messages generated during validation include whether any management classes or copy groups refer to storage pools that do not exist on the target server.
   - Query management class and copy group definitions to compare the storage destinations specified with the names of existing storage pools on the target server.

d. Depending on the amount of client file data that you expect to import, you may want to examine the storage hierarchy to ensure that sufficient storage space is available. The amount of space that is required can be deducted from the output of the preview process.

e. After you have defined the appropriate storage hierarchy on the target server, you can import file data from the tape volumes.

5. After completing the previous steps, the new Tivoli Storage Manager server is ready for operations. However, there is still the application server that needs to be modified. This step will be the same as the situation in

which we moved the library between equal operating systems (Section 6.2.3.1, “Moving the library to an equal operating system” on page 181). For details, refer to step 5 on page 183.

6.2.3.3 Using a new library

A final migration scenario is the case in which you are moving the server, using a new type of library instead of the one previously used.

> **Note**
>
> An important assumption is that the new library uses incompatible media compared to the one currently used. If this is not the case, use the previous scenario (see 6.2.3.2, “Moving the library between different server platforms” on page 184) as example. The only difference is that instead of physically moving the library, you will move the tape volumes containing the export data, and use them in the new library.

The main difference with the previous scenario in Section 6.2.3.2, “Moving the library between different server platforms” on page 184, is that the media is incompatible. This means that you cannot use the current local library as destination for the export media. The following solutions can be considered.

Using file device class

A first, and apparently straightforward solution would be to use a device class of the file type, which would generate sequential files on the server’s hard disk. These files can then be moved to the new server, and imported.

But, since you are exporting all your data, this operation is the same as copying all the data that resides in the tape storage pools back to disk. A common approach is to export/import only the data of one or a few clients at one time and repeat this process until all data are moved. This is a very time consuming process, and therefore, we will not go into further detail on this solution.

Using a virtual device class

The virtual device class allows you to use *virtual sequential volumes* on your Tivoli Storage Manager server, which physically reside on another Tivoli Storage Manager server. Using this technique, it would be possible to export the server to tapes residing in the library attached to the new server. The steps that need to be performed resemble those in the previous situation (Section 6.2.3.2, “Moving the library between different server platforms” on page 184), but they need to be performed in a different order.

Begin with preparing your original Tivoli Storage Manager server (see step 1 on page 184). Next, you should install your new server, adding the SAN library, and defining it to the server (see step 3 on page 186). Next set up the virtual volumes. This includes setting up server-to-server communications, and defining a client node on the target server, which has an archive copy group defined. Use the SAN library as target for this copy group. Finally, you should define a device class on the original server, which uses these virtual volumes.

Next, perform the export operation as described in step 2 on page 185, followed by the import and application server setup (see step 4 on page 186 and step 5 on page 187).

The major disadvantage of this solution is that you will transfer all your backup data over the LAN network, which might not be possible due to bandwidth restrictions.

Using a shared library

A final solution, which probably gives the best performance, is using your new SAN infrastructure to create the export media. If we look at Figure 87 on page 177, we see that all servers are connected to the SAN, using the new library. If in our migration process we would add one step, we could use the SAN library as shared library between the new Tivoli Storage Manager server, and the Tivoli Storage Manager server that will be migrated (see Figure 92). By doing this, the export process could use the media in this library.

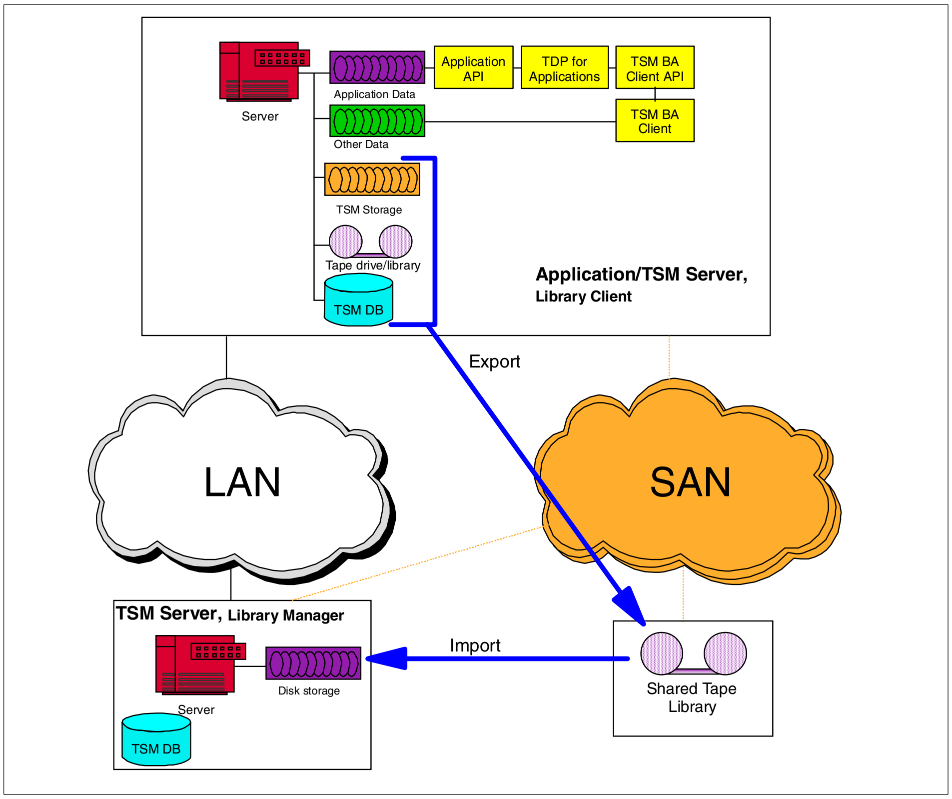

Figure 92. Sharing the library for export processing

Again, all steps are comparable to the ones described before, except for their order. Be careful to use the following order when performing this operation:

1. Prepare your original server for export processing, as described in step 1 on page 184.
2. Install the new Tivoli Storage Manager server, both hardware and software. Also define the required devices and storage needs, including the SAN attached library. See step 3 on page 186.
3. Next, install necessary hardware in the original machine, so that it is able to connect to the SAN shared library.
4. Set up library sharing between the original server and new server. The original server should be set as library client of the new Tivoli Storage Manager server, acting as library manager. For information on setting up library sharing, see Chapter 4, “Tape library sharing” on page 75.

5. Export the original server, as described in step 2 on page 185. Use the device class that points to the shared library.
6. After the export completes successfully, stop the original Tivoli Storage Manager server.
7. To finish the migration, perform the import (step 4 on page 186) and set up the Storage Agent on the application server (step 5 on page 187)

This completes the migration part. Next we, will look at two examples.

6.3 Examples of TSM migrations

This section will use the scenarios defined in 6.2, “Steps for migrating TSM into a SAN environment” on page 163 to migrate some sample configurations. Two examples are given:

- Consolidating servers, moving their libraries to a shared SAN library, together with a TDP for MS Exchange client backing up its data to one of these servers (see Section 6.3.1, “Shared library example” on page 191).
- Moving a local Tivoli Storage Manager server, combined with a TDP for MS Exchange client, to a new server, running a different operating system. The resulting configuration will use the Storage Agent software to back up the MS Exchange client (see Section 6.3.2, “LAN-free backup example” on page 209).

6.3.1 Shared library example

The first example that we will give is the migration of two Tivoli Storage Manager servers, each with its own SCSI attached library to a SAN environment, making use of a shared library. Additionally, one of the servers serves a client, using TDP for MS Exchange to make backups. This example will use the scenarios as described in Section 6.2.1, “Consolidation of multiple TSM server attached libraries” on page 165 and Section 6.2.2, “LAN attached application servers” on page 175.

The initial situation is shown in Figure 93. The first server is called Jamaica, and runs Windows 2000. Its storage consists of a disk storage pool migrating into a tape storage pool. The tape storage pool uses a SCSI attached Magstar 3570 library, with one drive. Its main job is to back up normal file clients, including himself.

The second server, Brazil, runs AIX 4.3.3. It has two disk storage pools, and two tape storage pools. Both of these pools reside on a Magstar 3570 MP library, also with one drive. Its job consists of backing up the application servers, running MS Exchange. This is done through the Tivoli Data Protection for MS Exchange server. An example of such a client is Diomede, running Windows NT 4.0. Besides backing up the application, it also stores the other files belonging to these application servers.

Note

Although both servers are using the same type of tape devices, we could use the exceptional case described in 6.2.1, “Consolidation of multiple TSM server attached libraries” on page 165, which states that when consolidating compatible libraries, it is enough to simply move the volumes. However, for this example we will follow the paths required for non-compatible libraries.

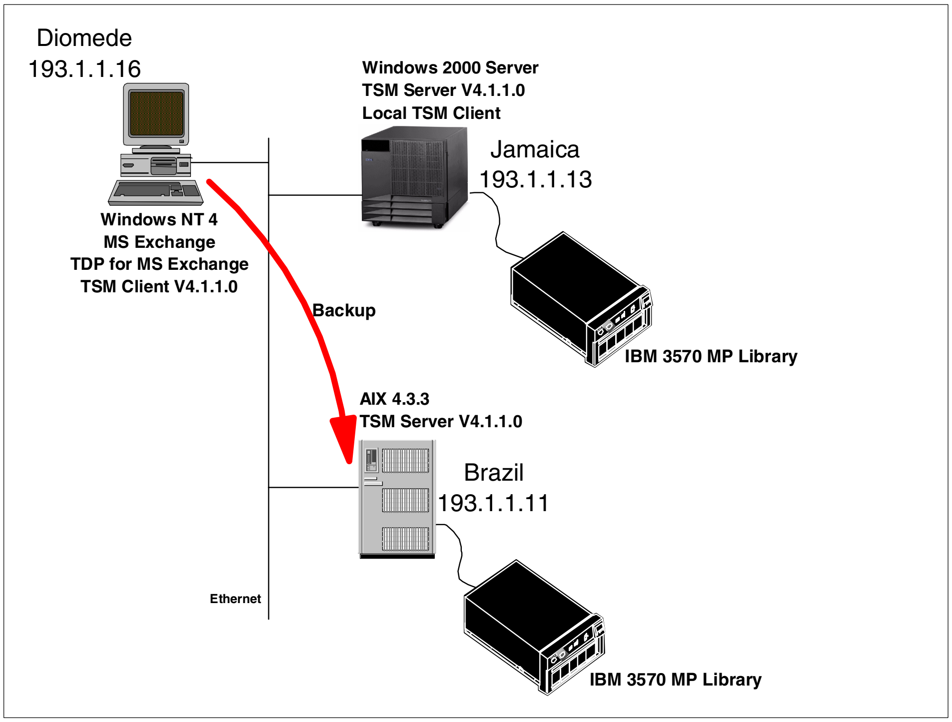

Figure 93. Initial configuration for shared library example

The goal of the migration, as shown in Figure 94, is to share the Magstar 3570 MP library between the two servers, and to use the Storage Agent component to provide a LAN-free solution for the TDP for MS Exchange client. The library manager will be Brazil, and the library client Jamaica.

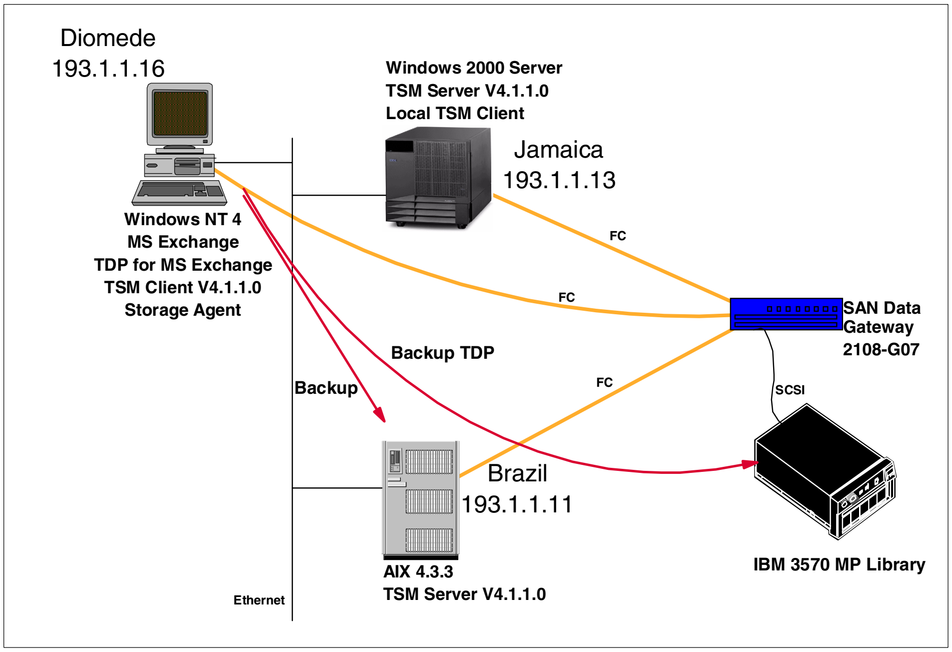

Figure 94. Example 1, goal of the migration

When we look at the start configuration, and the final goal, we can identify the following migration steps:

- Moving the SCSI attached library from the AIX system, to a SAN attached library on the same AIX server. The physical library used will remain the same.
- Adding the Windows 2000 server as library client to the AIX system, sharing the SAN attached library. This involves moving the storage pools using the original library.
- Installing a LAN-Free solution on the Windows NT application server.

The following sections will describe these actions in detail.

6.3.1.1 Moving the library on the AIX system

The first phase of this migration will be to attach the library belonging to the AIX machine to the SAN fabric. This move is shown in Figure 77 on page 167. The necessary hardware needs to be installed and configured on the RS/6000 machine. In addition, we must update the Tivoli Storage Manager configuration, and define Brazil as library server.

1. Installation of the necessary hardware.

   The first step is to attach the RS/6000 system and the library to the SAN fabric. We make the assumption that you have a SAN fabric installed. Also, the required device drivers will need to be installed, and Fibre Channel support added to the Tivoli Storage Manager server. For details on how to do this, please refer to Section 4.1.2.3, “Operating system preparation” on page 80. Once this is finished, the library is ready to be used by the Tivoli Storage manager.

2. Updating the Tivoli Storage Manager installation.

   Tivoli Storage Manager support for Fibre Channel protocol is delivered in the install package `tivoli.tsm.devices.fcprte`. This package and the SCSI-only support package, `tivoli.tsm.devices.rte`, are mutually exclusive. Since until now, the library was SCSI attached, you will need to remove the SCSI-only package, and install the FC support.

   To do this, use the **Software Maintenance and Utilities** menu item in SMIT, and select the `tivoli.tsm.devices.rte` file set. This also removes `tivoli.tsm.msg.en_US.devices`. When the remove is complete, return to the main **Software Installation and Maintenance** menu, and select the **SMIT Install and Update from ALL available software** menu. Select the following file sets for installation:

   - `tivoli.tsm.devices.fcprte`
   - `tivoli.tsm.msg.en_US.devices`

3. Setting up the shared library environment.

   After installing the required FC support for Tivoli Storage Manager, it is now time to define the AIX machine as library manager, and update the library definition, so that it can be shared.

   The first step, if not yet configured, is to set up the server so that it can use server-to-server communications, required to talk to the library clients. Enabling server-to-server communication means specifying the server’s name, password, and network addresses. This can be done by issuing the following commands:

```
set servername brazil
set serverpassword brazil
set serverhladdress 193.1.1.11
set serverlladdress 1500
```

These commands define the AIX server as Brazil, with password Brazil, using TCPIP address 193.1.1.11 and TCPIP port 1500. In addition, you can set the CROSSDEFINE parameter to *on*. Normally, *on* is the default value. This specifies that when a server defines another server for communications, the defining server is automatically defined on the target server.

4. After specifying the server's communication parameters, we must now redefine the 3570 library. The library in our case is called 3570LIB.

   Issue the `QUERY LIBRARY 3570LIB F=D` command. This will return the following information:

```
Session established with server BRAZIL: AIX-RS/6000
  Server Version 4, Release 1, Level 1.0
  Server date/time: 09/22/2000 13:58:21  Last access: 09/22/2000 12:47:31


                  Library Name: 3570LIB
                  Library Type: SCSI
                        Device: /dev/rmt3.smc
              Private Category:
              Scratch Category:
              External Manager:
                        Shared: No
       Primary Library Manager:
     Secondary Library Manager:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 09/20/2000 15:30:55
```

   First, update the device of the library and the tape drives in the library. To find the new device names, issue the following AIX command:

```
lsdev -Cc tape

rmt1 Available 10-68-01-2,0 IBM Magstar MP Tape Subsystem (FCP)
```

   Also, the shared status is set to No. This should be changed to Yes. Issue the following command to update the library and drive:

```
update library 3570lib device=/dev/rmt1.smc shared=yes
update drive 3570lib drive0 device=/deb/rmt1
```

Here are the results of these operations:

```
Query library
Session established with server BRAZIL: AIX-RS/6000
  Server Version 4, Release 1, Level 1.0
  Server date/time: 09/22/2000 13:58:21  Last access: 09/22/2000 12:47:31


                  Library Name: 3570LIB
                  Library Type: SCSI
                        Device: /dev/rmt1.smc
              Private Category:
              Scratch Category:
              External Manager:
                        Shared: Yes
       Primary Library Manager:
     Secondary Library Manager:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 09/20/2000 15:30:55

Query drive

                                Library Name: 3570LIB
                                  Drive Name: DRIVE0
                                 Device Type: 3570
                                      Device: /dev/rmt1
                                     ON LINE: Yes
                                     Element: 16
                                Allocated to:
              Last Update by (administrator): ADMIN
                       Last Update Date/Time: 10/04/2000 11:07:10
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE
```

This completes the required setup on the AIX system. By defining the library as shared, Brazil is implicitly defined as library manager for that library. All other library users or clients will have to point to Brazil as library manager in order to access the library.

Since no other definitions have changed, the device class, storage pools, and storage pool volumes associated with the library still exist.

6.3.1.2 Moving the library on the Windows 2000 system

The second step in our migration will be moving the library storage pools from the local library attached to the Windows 2000 system to the shared library. As stated before, we will assume that both libraries use incompatible media. Before we start moving, it is a good idea to look at the storage configuration of Jamaica. The main job will be to identify the storage pools, and how they are used. Figure 95 shows the layout of the storage environment of the Windows 2000 system.

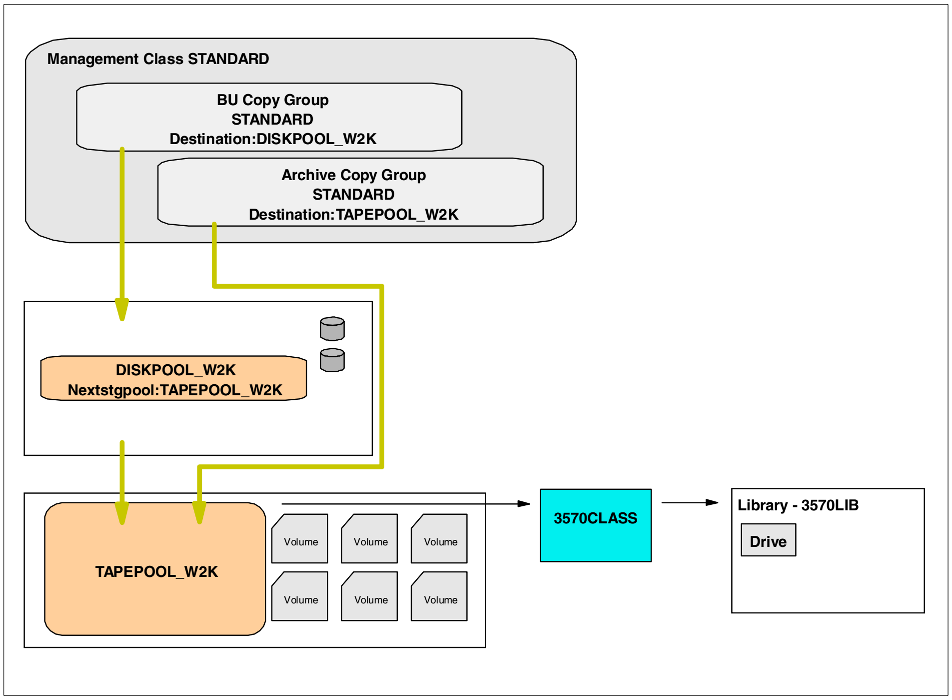

Figure 95. Storage definitions, initial situation

Two storage pools exists, one on disk, and one using the tape library. The tape pool, TAPEPOOL_W2K, will use the library, 3570LIB_W2k, using a device class 3570CLASS. As stated in Section 6.2.1, “Consolidation of multiple TSM server attached libraries” on page 165, you should establish which storage pools will be moved. In our case, there is only one storage pool that will be migrated, TAPEPOOL_W2K. The disk pool will remain local on the server.

This layout means that we will have to follow the scenarios described in Section 6.2.1.1, “Migration-receiving storage pools” on page 168 and Section 6.2.1.2, “Directly node-attached sequential storage pools” on page 173. The following sections will show how to do this in detail.

Preparing the server

The first step is to prepare the server in such a way that we will be able to recover to the initial situation if there is a problem, and disable all actions that might interfere with our migration process.

Start by performing a database backup. In this case, we will back up the database to a volume in the local 3570 library, 3570LIB belonging to the 3570CLASS device class. To do this, use the following command:

```
Backup database devclass=3570class type=dbsnapshot
```

This command will create a database backup, using a scratch volume from the 3570 library.

Next, disable the server for client sessions, using the following command:

```
Disable Sessions
```

Also, you might want to block general administrator access, only keeping the one you are working with. This can be done using the `LOCK ADMIN` command. Also, check all administrative schedules that might run and interfere with the migration process. You can de-activate these schedules.

After this is finished, it is now necessary to prepare the storage pool that we are migrating, including all other storage pools in the hierarchy that might be using him. In our case, this means updating DISKPOOL_W2K and TAPEPOOL_W2K. The first step will be to ensure that the disk storage pool does not move its data to the tape storage pool, using a migration process. To block this, issue the following update:

```
Update stgpool DISKPOOL_W2K highmig=100
```

For the tape pool, we should place all volumes in read/only mode, protecting them against any accidental modification. Do this using the following statement:

```
Update volume * access=readonly wherestgpool=TAPEPOOL_W2K
```

Next, update the tape storage pool so that it does not request new volumes, and that it keeps the volumes you will move (instead of returning them to scratch). This will enable you to recover easily in the event of an error:

```
Update stgpool TAPEPOOL_W2K maxscratch=0 reusedelay=10
```

This completes the server preparation. After this, halt the server and start installing the necessary hardware and drivers for SAN attachment.

Attaching the server to the SAN

If the current server has no Fibre Channel attachment, it is now time to do this. This step will involve installing one or more Fibre Channel host bus adapters or HBAs in the system, and connecting it to the SAN. For Netfinity servers, the Qlogic 2100 and 2200 series Fibre Channel adapters are supported. Firmware and drivers for these adapters are available from the following URL:

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/NT/sdgnt.htm`

For details on installing and configuring these drivers, refer to Section 4.1.2.3, “Operating system preparation” on page 80.

Defining the shared library

After installing the necessary hardware and drivers for the HBA, and having connected this system to the SAN, we can now start defining the new shared library. Since this library will be shared between this system and the AIX server Brazil, which is also the library manager, ensure that the AIX server is running and has access to the library. Also, check if Windows 2000 sees the SAN attached library and drives. This can be done using the Tivoli Storage Manager Server utilities, and looking at the devices shown.

In our example, we see the SAN library and drive, with the following properties:

Library	lb1.3.0.3
Drive 0	mt1.2.0.3

The first step will be to define the shared library, using the `define library` command. Although this is a very simple command, there is one problem in our environment. A requirement for the library name in a shared environment is that all library clients, like Windows 2000 system, must use the exact same name for the library as it is defined on the library manager. Referring to Section 6.3.1.1, “Moving the library on the AIX system” on page 193, we named the library 3570LIB. So, the library definition on the Windows 2000 system must be done using the same name. But, since the local library on the Windows 2000 system also uses this name, we cannot perform this definition. To solve this problem, we must first rename the local library, for example, using 3570LIB_W2K.

Start by determining the parameters used for the current library and drive definition:

```
Query library 3570LIB f=d
Query drive drive0 f=d
```

The output is as follows:

```
Query Library

                  Library Name: 3570LIB
                  Library Type: SCSI
                        Device: LB3.1.0.1
              Private Category:
              Scratch Category:
              External Manager:
                RSM Media Type:
                        Shared: No
       Primary Library Manager:
     Secondary Library Manager:
Last Update by (administrator): ADMIN
         Last Update Date/Time: 10/04/2000 10:08:46

Query Drive


                                Library Name: 3570LIB
                                  Drive Name: DRIVE0
                                 Device Type: 3570
                                      Device: MT3.0.0.1
                                     ON LINE: Yes
                                     Element: 17
                                Allocated to:
              Last Update by (administrator): ADMIN
                       Last Update Date/Time: 10/04/2000 10:09:10
Cleaning Frequency (Gigabytes/ASNEEDED/NONE): NONE
```

Define a new library and drive, using the same device name and library type:

```
Define library 3570LIB_W2K libtype=SCSI device=lb3.1.0.1
Define drive 3570LIB_W2K drive0_w2k device=mt3.0.0.1 element=17
```

This completes the physical definition of the local library. Next, update your device class so that it uses this new library name:

```
Update devclass 3570CLASS devtype=3570 library=3570LIB_W2K
```

Finally, the tape volumes that reside physically in the library should be logically checked in using the new library name. Start by checking them out of the 3570LIB, ensuring that you do not remove them.

```
checkout libvol 3570lib 081df3 checklabel=no remove=no
```

Repeat this command for all volumes inside the library. Next, check the same volumes in, using the new library name:

```
Checkin libvolume 3570LIB_W2K search=yes status=private
```

As you can see, we check the volumes in as private. This is enough, since no new scratch volumes will be required.

Finally, delete the original drive and library.

```
delete drive 3570lib drive0
delete library 3570lib
```

This completes the rename of the original library. After this rename, we can start defining the shared library. Since library sharing is going to use server-to-server communications, start by enabling the Windows 2000 server for this purpose, and define the AIX server.

```
set servername jamaica
set serverpassword jamaica
set serverhladdress 193.1.1.11
set serverlladdress 1500
set crossdefine on
Define server brazil serverpa=brazil hla=193.1.1.11 lla=1500 crossd=yes
```

To check if the server definitions are OK, issue the following command on both Jamaica and Brazil:

```
Query server f=d
```

On Jamaica, this command results in the following:

```
                         Server Name: BRAZIL
                        Comm. Method: TCPIP
                  High-level Address: 193.1.1.11
                   Low-level Address: 1500
                         Description:
                   Allow Replacement: No
                           Node Name:
               Last Access Date/Time: 10/04/2000 12:05:27
              Days Since Last Access: <1
                            Locked?: No
                         Compression: No
             Archive Delete Allowed?: (?)
                                 URL:
              Registration Date/Time: 10/04/2000 12:05:27
            Registering Administrator: ADMIN
          Bytes Received Last Session: 0
              Bytes Sent Last Session: 0
             Duration of Last Session: 0.00
          Pct. Idle Wait Last Session: 0.00
         Pct. Comm. Wait Last Session: 0.00
         Pct. Media Wait Last Session: 0.00
               Grace Deletion Period: 5
                    Managing profile:
                 Server Password Set: Yes
        Server Password Set Date/Time: 10/04/2000 12:05:27
        Days Since Server Password Set: <1
      Invalid Sign-on Count for Server: 0
          Virtual Volume Password Set: No
  Virtual Volume Password Set Date/Time: (?)
 Days Since Virtual Volume Password Set: (?)
Invalid Sign-on Count for Virtual Volume Node: 0
```

On Brazil, the following server definitions are present:

```
                              Server Name: JAMAICA
                             Comm. Method: TCPIP
                       High-level Address: 193.1.1.11
                        Low-level Address: 1500
                              Description:
                        Allow Replacement: No
                                Node Name: BRAZIL
                   Last Access Date/Time: 10/04/2000 12:05:36
                  Days Since Last Access: <1
                                  Locked?: No
                              Compression: No
                  Archive Delete Allowed?: (?)
                                      URL:
                  Registration Date/Time: 10/04/2000 12:05:36
               Registering Administrator: JAMAICA
             Bytes Received Last Session: 0
                 Bytes Sent Last Session: 0
                Duration of Last Session: 0.00
             Pct. Idle Wait Last Session: 0.00
            Pct. Comm. Wait Last Session: 0.00
            Pct. Media Wait Last Session: 0.00
                   Grace Deletion Period: 5
                        Managing profile:
                     Server Password Set: Yes
           Server Password Set Date/Time: 10/04/2000 12:05:36
          Days Since Server Password Set: <1
         Invalid Sign-on Count for Server: 0
             Virtual Volume Password Set: Yes
   Virtual Volume Password Set Date/Time: 10/04/2000 12:05:36
  Days Since Virtual Volume Password Set: <1
Invalid Sign-on Count for Virtual Volume Node: 0
```

Figure 96 shows the definitions we must make for the new, shared library and attached storage pools.

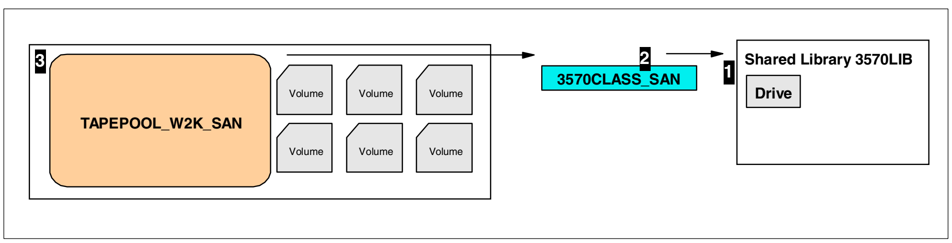

Figure 96. Defining the shared library

The following commands must be performed:

```
define library 3570LIB libtype=shared primarylibmanager=brazil
define drive 3570LIB drive0 device=mt1.2.0.3
define devclass 3570CLASS_SAN devtype=3570 library=3570LIB
define stgpool TAPEPOOL_W2K_SAN 3570CLASS_SAN maxscratch=5
```

For these definitions, use the current configuration to determine other parameters that are required, like the MAXSCRATCH value, reclamation values, collocation, next storage pools, the MAXSIZE, reuse and migration delays, drive format, MOUNTRETENTION and MOUNTWAIT times, and the mount limit.

After completing this step, we are actually going to move the data from the local library to the shared library at the SAN, as shown in Figure 97.

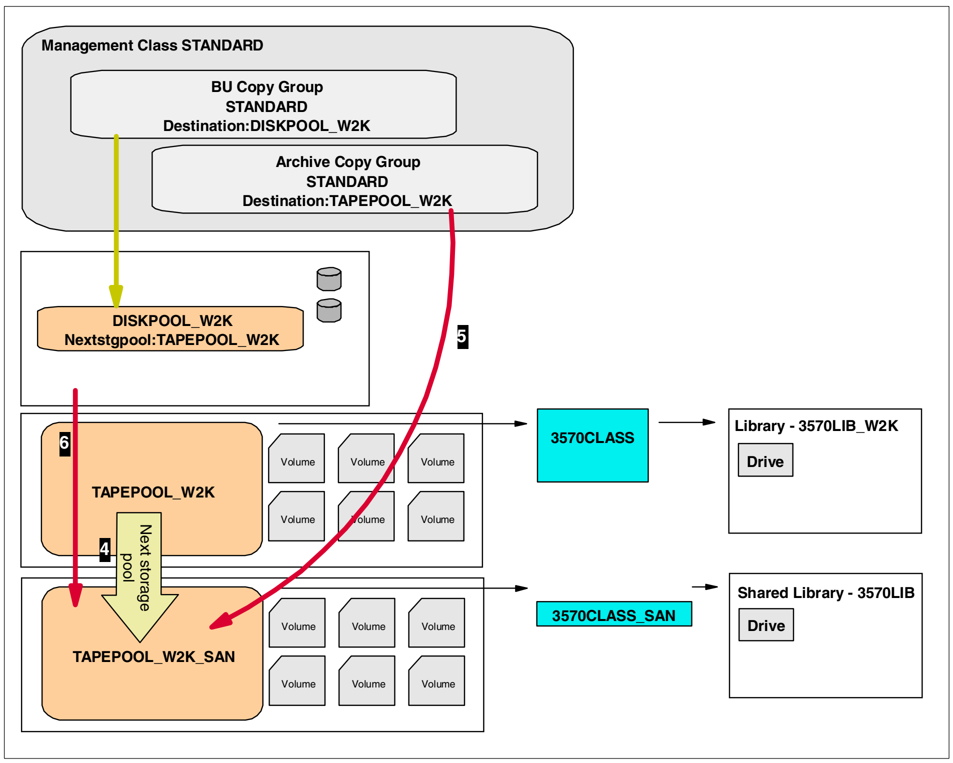

Figure 97. Redefining storage

Before you can start, check that the library contains the necessary amount of scratch volumes to accommodate the original storage pool. If not, the check-in of scratch volumes should be done on the library manager. To determine this number, look at the current number of volumes assigned to the storage pool TAPEPOOL_W2K. Keep in mind that moving the data will consolidate volumes, filling them as high as possible, unless collocation was set on the new TAPEPOOL_W2K_SAN pool.

Start by moving the data from the current storage pool to the new shared one. We will do this by defining TAPEPOOL_W2K_SAN as next storage pool for TAPEPOOL_W2K and invoking the migration process:

```
Update stgpool TAPEPOOL_W2K nextstgpool=TAPEPOOL_W2K_SAN
Update stgpool TAPEPOOL_W2K lowmig=0 highmig=0
```

This command will start a migration process, as follows:

```
 Process Process Description  Status
  Number
-------- -------------------- -------------------------------------------------
       2 Migration            Volume 081DF3 (storage pool TAPEPOOL_W2K), Moved
                               Files: 4301, Moved Bytes: 421,250,168,
                               Unreadable Files: 0, Unreadable Bytes: 0.
                               Current Physical File (bytes): 5,305,854

                               Current input volume: 081DF3.

                               Current output volume: 08644C.
```

When the migration process completes successfully, modify the definitions so that you will use the new shared pool instead of the old one. Looking at Figure 95 on page 197, we see that these modifications must be done in two places:

- The disk storage pool
- The archive copy group

```
Update stgpool DISKPOOL_W2K nextstgpool=TAPEPOOL_W2K_SAN
Update copygroup STANDARD STANDARD STANDARD STANDARD type=archive destination=TAPEPOOL_
```

After issuing these commands, we must commit the changes by activating the policy set in which we modify the copy group and re-enable the server for a client session, using the following commands:

```
Activate policyset STANDARD STANDARD
Enable sessions
```

Cleanup and final modifications

This step completes the shared library implementation for client backups. What should still be done is to clean up the original definitions, and ensure that all objects referring to the original storage pool and device class have been modified.

Some objects that might still refer to TAPEPOOL_W2K are administrative schedules, like migration and reclamation. Objects that might refer to the original device class are database backups and storage pool backups. The latter ones will also use a copy storage pool, which might have been assigned to the original device class. Delete the original copy storage pools, and redefine them on the new shared library.

The other actions are to delete the original storage pool and the volumes assigned to it, delete the device class, and finally the library and drives. If you wish to keep these objects for a while, while ensuring that the new configuration works, you can. But, since you have the database backup, you will be able to restore the situation if you keep the tape volumes (physically).

6.3.1.3 Set up the LAN-free TDP for MS Exchange backup

As the last step in our migration, we will set up the application server Diomede so that it is able to back up the MS Exchange databases directly over the SAN. This example is described in Section 6.2.2, “LAN attached application servers” on page 175.

Figure 98 shows the layout of the storage pools and management classes defined on Brazil, which is the server Diomede uses as a backup destination (see Figure 93 on page 192).

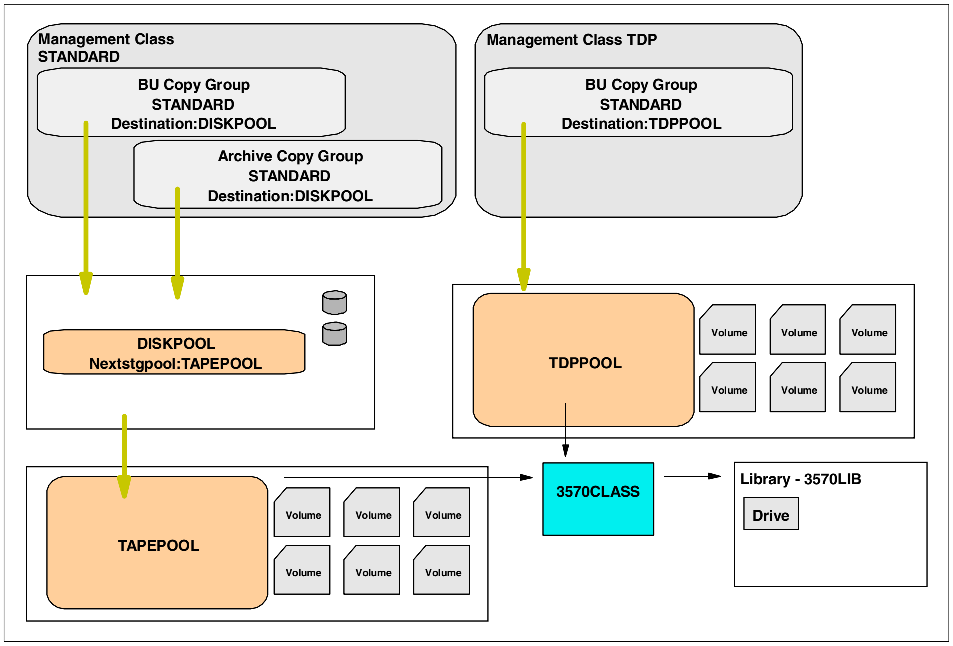

Figure 98. Storage setup for the AIX server

The data movement between Diomede and Brazil is done entirely over the LAN. The goal of this migration is to get the TDP data moved over the SAN instead.

Since the TDP client used the TDP management class, it sends its data directly to tape. Therefore, we do not need to move data from a disk storage pool to this tape pool (as stated in step 2 on page 178).

The following steps were performed during this migration:

1. Install the necessary hardware and drivers to connect Diomede to the SAN. Please refer to Section 5.2.2, “Verify configuration of SAN devices on the client” on page 130 for details on how to do this.
2. Install the Storage Agent software, as described in Section 5.3.1, “Installing the Storage Agent Software” on page 133.
3. Check the DSMTA.OPT file in the Storage Agent directory for the following required statements:

```
COMMETHOD TCPIP
COMMETHOD NAMEDPIPES
DEVCONFIG devconfig.txt
```

4. Start the ADSMSCSI driver in the devices application. After starting the the device driver, use the `TSMDLST.EXE` application to determine the SCSI ID of the tape in our shared library.

```
Computer Name:      DIOMEDE
TSM Device Driver:  Running

TSM Device Name  ID   LUN  Bus  Port   TSM Device Type  Device Identifier

--------------------------------------------------------------------------------
mt0.2.0.2        0    2    0    2      3570             IBM     03570C125424
lb0.3.0.2        0    3    0    2      LIBRARY          IBM     03570C125424
```

5. With this information, we can now configure Brazil, so that it will allow the Storage Agent to work directly with the drive in the shared library. This configuration includes defining a server on Brazil, and a drive mapping. The following commands were used:

```
define server sa_diomede serverpassword=dio hla=193.1.1.16 lla=1500
define drivemapping sa_diomede 3570lib drive0 device=mt0.2.0.2
```

When defining the Storage Agent on Brazil, do not use the cross define function.

6. Next, the Storage Agent must be configured so that it can connect to Brazil for drive sharing. This is done using the following command:

```
dsmsta setstorageserver myname=sa_diomede mypassword=dio servername=brazil
serverpassword=brazil hladdress=193.1.1.11 l laddress=1500
```

This command will generate the a device configuration file for the Storage Agent, called `devconfig.txt.`

```
SET STANAME SA_DIOMEDE
SET STAPASSWORD 18dd2370
DEFINE SERVER BRAZIL HLADDRESS=193.1.1.11 LLADDRESS=1501 SERVERPA=185b801eb97c22
```

7. Finally, modify the option file that will be used by the TDP for MS Exchange client (DSM.OPT), so that it will use the Storage Agent rather than the Tivoli Storage Manager server (Brazil) to perform backups and restores. Add the following line to the option file:

```
ENABLELANFREE YES
```

8. Once this has been done, you can start the Storage Agent. It is advised to run this agent as a service. Refer to 5.3.6, “Installing the Storage Agent as a service” on page 139 for details.

Important

When using the Storage Agent, the filespaces on the server for the MS Exchange client will not change. This means that you can see MS Exchange backups created before the migration (using a LAN connection). However, if you try to restore these backups using the Storage Agent, the transaction will fail. In order to solve this, you should stop the Storage Agent service whenever restoring a pre-SAN backup. By doing so, you will reconnect to the Tivoli Storage Manager server using the LAN connection, and the restore will complete.

This completes our first migration example, in which we moved a library from a locally attached one to a shared SAN library. Next, we moved the data from another server to this shared library, making him a library client. Finally, we implemented the Storage Agent on a system running the TDP for MS Exchange application.

6.3.2 LAN-free backup example

A second migration example will focus on consolidating a stand-alone backup solution using Tivoli Storage Manager server on one machine with another Tivoli Storage Manager server using LAN-free backup in a SAN environment. This method will be necessary when you are currently using a local Tivoli Storage Manager server on your application server, and want to move to a SAN attached shared library, using the Storage Agent. Details about this type of migration are given in Section 6.2.3, “Locally attached tape library” on page 180.

Figure 99 shows the initial layout for this example.

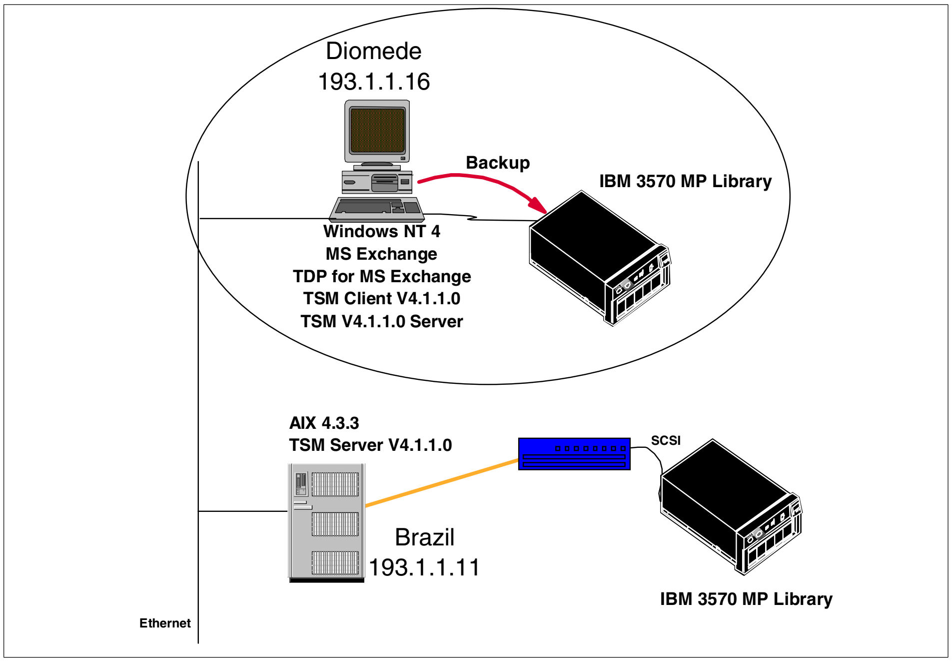

Figure 99. Initial layout for LAN-free backup example

Diomede, our application server, runs a local Tivoli Storage Manager server to back up the MS Exchange data to a locally attached SCSI tape library sharing. Brazil, a Tivoli Storage Manager Server running on AIX, has a shared, SAN connected tape library. In the starting situation, however, both servers do not communicate. You could say that Diomede lives on an island, needing extra management and operations.

The goal of the migration, as shown in Figure 100, is to have Diomede back up its MS Exchange data to the shared library using the SAN, and its file data to Brazil, using the normal network. The main focus, however, is to perform this move, with the least possible impact on current backup data for Diomede. This means that after the move, the data that is currently residing locally should be still be available.

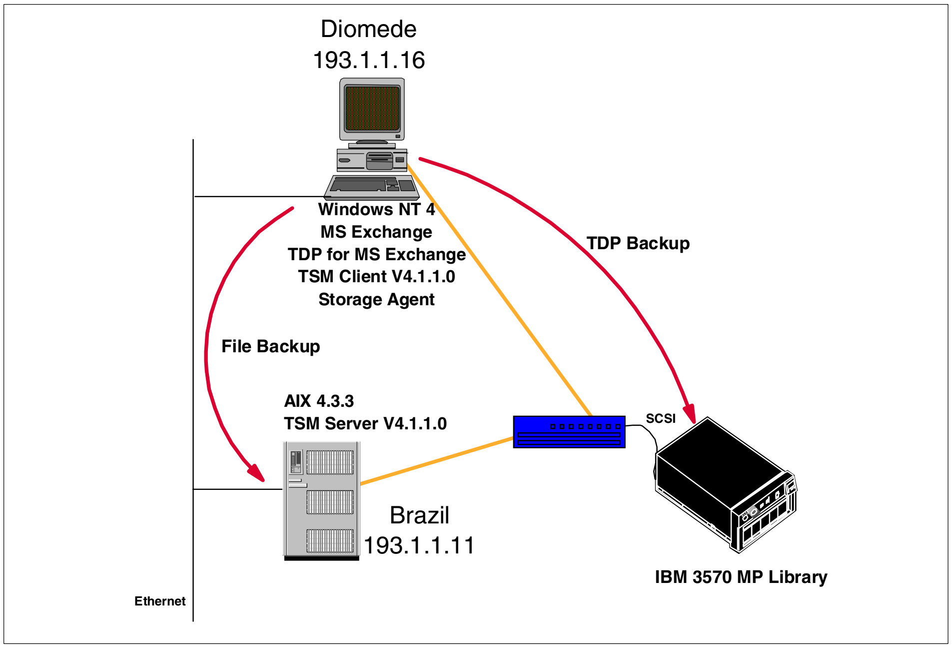

Figure 100. Goal of LAN-free backup example

The following steps were performed to do this migration (for details, see Section 6.2.3.2, “Moving the library between different server platforms” on page 184).

1. Back up the Tivoli Storage Manager server database on Diomede:

```
Backup database devclass=3570class type=dbsnapshot
```

2. Disable all client sessions:

```
Disable sessions
```

3. The next step would be to move the data from the disk storage pool used by the TDP client to the tape storage pool. This due to the fact that after the migration, the Storage Agent will require that you only use a tape storage pool. In our example, however, as shown in Figure 101, this is not required, since no disk storage pool is used for the TDP client (TDP management class).

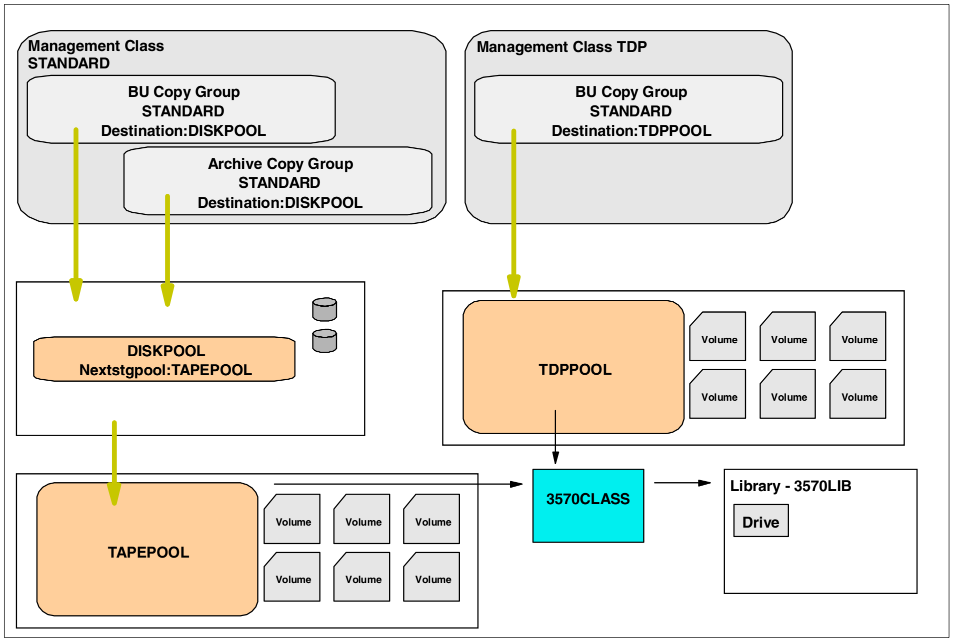

Figure 101. Original storage setup on Diomede

If you find that you need to do this, use the following commands:

```
Move data diskvolume_name stgpool=TDPTAPEPOOL
Update copygroup TDP TDP TDP STANDARD destination=TDPTAPEPOOL
Activate policyset TDP TDP
Delete copygroup diskpool_name
```

4. Next, gather information on your current system that might be required for recovery in case of failure, and to define the necessary storage pools on Brazil.

```
Backup devconfig filenames=C:\devconfig.out
Backup volhist filenames=C:\volhist.out
```

```
q stgpool f=d
               Storage Pool Name: DISKPOOL
               Storage Pool Type: Primary
               Device Class Name: DISK
         Estimated Capacity (MB): 24.0
                        Pct Util: 0.0
                        Pct Migr: 0.0
                     Pct Logical: 100.0
                    High Mig Pct: 90
                     Low Mig Pct: 70
                 Migration Delay: 0
              Migration Continue: Yes
             Migration Processes: 1
               Next Storage Pool: TAPEPOOL
            Reclaim Storage Pool:
          Maximum Size Threshold: No Limit
                          Access: Read/Write
                     Description:
               Overflow Location:
           Cache Migrated Files?: No
                      Collocate?:
           Reclamation Threshold:
 Maximum Scratch Volumes Allowed:
   Delay Period for Volume Reuse:
          Migration in Progress?: No
            Amount Migrated (MB): 23.58
Elapsed Migration Time (seconds): 30
        Reclamation in Progress?:
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
           Last Update Date/Time: 10/04/2000 19:02:08

               Storage Pool Name: TAPEPOOL
               Storage Pool Type: Primary
               Device Class Name: 3570CLASS
         Estimated Capacity (MB): 5,000.0
                        Pct Util: 1.8
                        Pct Migr: 100.0
                     Pct Logical: 100.0
                    High Mig Pct: 90
                     Low Mig Pct: 70
                 Migration Delay: 0
              Migration Continue: Yes
             Migration Processes:
               Next Storage Pool:
            Reclaim Storage Pool:
          Maximum Size Threshold: No Limit
                          Access: Read/Write
                     Description:
               Overflow Location:
           Cache Migrated Files?:
                      Collocate?: No
           Reclamation Threshold: 60
 Maximum Scratch Volumes Allowed: 1
   Delay Period for Volume Reuse: 0 Day(s)
          Migration in Progress?: No
            Amount Migrated (MB): 0.00
```

```
Continued from previous page ..

            Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
        Reclamation in Progress?: No
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
           Last Update Date/Time: 10/04/2000 19:01:34

               Storage Pool Name: TDPPOOL
               Storage Pool Type: Primary
               Device Class Name: 3570CLASS
         Estimated Capacity (MB): 5,000.0
                        Pct Util: 0.0
                        Pct Migr: 100.0
                     Pct Logical: 100.0
                    High Mig Pct: 90
                     Low Mig Pct: 70
                 Migration Delay: 0
              Migration Continue: Yes
             Migration Processes:
               Next Storage Pool:
            Reclaim Storage Pool:
          Maximum Size Threshold: No Limit
                          Access: Read/Write
                     Description:
               Overflow Location:
           Cache Migrated Files?:
                      Collocate?: No
           Reclamation Threshold: 60
 Maximum Scratch Volumes Allowed: 1
   Delay Period for Volume Reuse: 0 Day(s)
          Migration in Progress?: No
            Amount Migrated (MB): 0.00
Elapsed Migration Time (seconds): 0
        Reclamation in Progress?: No
 Volume Being Migrated/Reclaimed:
  Last Update by (administrator): ADMIN
           Last Update Date/Time: 10/04/2000 19:01:54
```

5. When we have this information, we can now start the export process. The first step will be to run the export in preview mode. This will give us an idea of what is exported, and how much storage space will be required. To perform the preview, run the following command:

```
export server filedata=all preview=yes
```

This will start a process, which will redirect its output to the server's activity log. The output in this log will look like this:

```
ANR0609I EXPORT SERVER started as process 8.
ANR0610I EXPORT SERVER started by ADMIN as process 8.
ANR0639I EXPORT SERVER: Processing domain STANDARD.
ANR0639I EXPORT SERVER: Processing domain TDP.
ANR0640I EXPORT SERVER: Processing policy set ACTIVE in policy domain STANDARD.
ANR0640I EXPORT SERVER: Processing policy set STANDARD in policy domain STANDARD.
ANR0640I EXPORT SERVER: Processing policy set ACTIVE in policy domain TDP.
ANR0640I EXPORT SERVER: Processing policy set TDP in policy domain TDP.
ANR0641I EXPORT SERVER: Processing management class STANDARD in domain STANDARD,
set ACTIVE.
ANR0641I EXPORT SERVER: Processing management class STANDARD in domain STANDARD,
set STANDARD.
ANR0641I EXPORT SERVER: Processing management class TDP in domain TDP, set ACTIVE.
ANR0641I EXPORT SERVER: Processing management class TDP in domain TDP, set TDP.
ANR0643I EXPORT SERVER: Processing archive copy group in domain STANDARD,
set ACTIVE, management class STANDARD.
ANR0643I EXPORT SERVER: Processing archive copy group in domain STANDARD,
set STANDARD, management class STANDARD.
ANR0642I EXPORT SERVER: Processing backup copy group in domain STANDARD,
set ACTIVE, management class STANDARD.
ANR0642I EXPORT SERVER: Processing backup copy group in domain STANDARD,
set STANDARD, management class STANDARD.
ANR0642I EXPORT SERVER: Processing backup copy group in domain TDP,
set ACTIVE, management class TDP.
ANR0642I EXPORT SERVER: Processing backup copy group in domain TDP, set TDP,
management class TDP.
ANR0638I EXPORT SERVER: Processing administrator ADMIN.
ANR0638I EXPORT SERVER: Processing administrator CLIENT.
ANR0638I EXPORT SERVER: Processing administrator DIOMEDE.
ANR0638I EXPORT SERVER: Processing administrator TDP_DIOMEDE.
ANR0635I EXPORT SERVER: Processing node CLIENT in domain STANDARD.
ANR0635I EXPORT SERVER: Processing node DIOMEDE in domain STANDARD.
ANR0637I EXPORT SERVER: Processing file space \\diomede\c$ for node DIOMEDE.
ANR0635I EXPORT SERVER: Processing node TDP_DIOMEDE in domain TDP.
ANR0637I EXPORT SERVER: Processing file space IBM.DOMAIN.DIOMEDE.DIR for node
TDP_DIOMEDE.
ANR0645I EXPORT SERVER: Processing schedule DAILY_INCR in domain STANDARD.
ANR0645I EXPORT SERVER: Processing schedule WEEKLY_INCR in domain STANDARD.
ANR2034E QUERY SCHEDULE: No match found using this criteria.
ANR0892I EXPORT SERVER: No matching optionsets found for exporting.
ANR0616I EXPORT SERVER: Preview processing completed successfully.
ANR0620I EXPORT SERVER: Copied 2 domain(s).
ANR0621I EXPORT SERVER: Copied 4 policy sets.
ANR0622I EXPORT SERVER: Copied 4 management classes.
ANR0623I EXPORT SERVER: Copied 6 copy groups.
ANR0624I EXPORT SERVER: Copied 2 schedules.
ANR0625I EXPORT SERVER: Copied 4 administrators.
ANR0626I EXPORT SERVER: Copied 3 node definitions.
ANR0627I EXPORT SERVER: Copied 2 file space 0 archive files, 195 backup files, and 0
space managed files.
ANR0629I EXPORT SERVER: Copied 92432381 bytes of data.
```

This information tells us that this export operation would export 88MB of data (message ANR0629I). It also shows which objects could be exported, and which will be exported. These include:

- Policy domains, 2 exported in our case.
- Policy sets, 4 in our case.
- Management classes, 4 in our case.
- Copy groups, 6 in our case.
- Administrative and client schedules, 2 in our case.
- Node definitions, 3 in our case
- File spaces
- Administrators
- Option sets

6. The next step will be to actually perform the export we previewed. To do this, you will need to specify a device class pointing to a sequential device. The best solution in our case would be to export the data to tape, which can immediately be reused on the AIX side. The problem, however, is that we have only one tape device, and that an export process requires at least two drives. The reason for this is that the process needs to read data from tape (the filespaces), and write to tape (the export data). Solutions for this type of problem are given in Section 6.2.3.3, “Using a new library” on page 188.

   In our example, we opted for a file device class, called EXPORTCLASS. This due to fact that the amount of data in our example was rather small. The following shows the command that was issued:

```
export server filedata=all preview=no devclass=exportclass scratch=yes
```

   Again, this command starts a process, which can be monitored by the QUERY PROCESS command:

```
 Process Process Description  Status
  Number
-------- -------------------- -------------------------------------------------
      10 EXPORT SERVER        ANR0648I Have copied the following: 2 Domains  4
                               Policy Sets  2 Schedules  4 Management Classes
                               6 Copy Groups  4 Administrators  3 Nodes  2
                               Filespaces  133 Backup Files  69611443 Bytes  (0
                               errors have been detected).

                               Current input volume: 081DF3.

                               Current output volume: D:\EXPORT\70713154.EXP.
```

When the process completes, the result of the process can be reviewed in the Tivoli Storage Manager server activity log. The output resembles the one that was created during the preview run, but adds information on which volumes that were used for the export operation.

7. After the export, we must now import the data on Brazil. First, you should move the volumes containing the export data so that Brazil can access them. In case the export media are on tape, checking them in on the library used by Brazil would be sufficient. In our case, however, we moved the files to Brazil’s file system, and created a file device class that points to this file system. You will also need to provide the input volumes for the import process. If there are many volumes, you can create an ASCII file which contains these volume names. For each volume that you will use, add a line with the volume name.

   When the volumes are available, run the import process in preview mode:

   ```
   import server filedata=all preview=yes devclass=importclass
   volumenames=FILE:/tmp/files.lst
   ```

   The result of this process will be available in the Tivoli Storage Manager server activity log.

8. Next, check the list of original storage pools, and define them on Brazil if they do not exist. In our case, the only storage pool that was not available on Brazil was TDPTAPEPOOL.

   ```
   Define stgpool TDPPOOL 3570CLASS maxscratch=1
   ```

   If you do not define this storage pool prior to the import, the process will fail with the following message:

   ```
   ANR0665W IMPORT SERVER: Transaction failed - storage pool TDPPOOL is not
   defined.
   ```

9. If the storage pools are defined, you can start the import processing. By default, none of the definitions currently on the server will be replaced. If you which that it does, add the REPLACEDEFS=YES parameter to the command.

```
import server filedata=all preview=no devclass=importclass
volumenames=FILE:/tmp/files.lst
```

10.When the import process completes successfully, you still need to set up Brazil and Diomede in such a way that Diomede acts as a backup client to Brazil for file data, and as a Storage Agent for MS Exchange backups.

Start by de-installing and removing the Tivoli Storage Manager server code, including the device driver currently installed on Diomede. This is required due to the fact that the Storage Agent cannot coexist with the Tivoli Storage Manager server.

11.Modify the Backup/Archive client option file so that the TCPSERVERADDRESS points to Brazil instead of Diomede.

12.Next, the Storage Agent must be installed and configured, and Brazil must be configured so that it shares the library with Diomede. Start by installing the necessary hardware and drivers to connect Diomede to the SAN. Please refer to Section 5.2.2, “Verify configuration of SAN devices on the client” on page 130 for details on how to do this.

13.Install the Storage Agent software, as described in Section 5.3.1, “Installing the Storage Agent Software” on page 133.

14.Check the DSMSTA.OPT file in the Storage Agent directory for the following required statements:

```
COMMETHOD TCPIP
COMMETHOD NAMEDPIPES
DEVCONFIG devconfig.txt
```

15.Start the ADSMSCSI driver in the devices application. After starting the the device driver, use the `TSMDLST.EXE` application to determine the SCSI ID of the tape in our shared library.

```
Computer Name:      DIOMEDE
TSM Device Driver:  Running

TSM Device Name  ID   LUN  Bus  Port   TSM Device Type  Device Identifier

-------------------------------------------------------------------------------------
mt0.2.0.2        0    2    0    2      3570             IBM     03570C125424
lb0.3.0.2        0    3    0    2      LIBRARY          IBM     03570C125424
```

16.With this information, we can now configure Brazil, so that it will allow the Storage Agent to work directly with the drive in the shared library. This configuration includes defining a server on Brazil, and a drive mapping. The following commands were used:

```
define server sa_diomede serverpassword=dio hla=193.1.1.16 lla=1500
define drivemapping sa_diomede 3570lib drive0 device=mt0.2.0.2
```

When defining the Storage Agent on Brazil, do not use the cross define function.

17.Next, the Storage Agent must be configured so that it can connect to Brazil for drive sharing. This is done using the following command:

```
dsmsta setstorageserver myname=sa_diomede mypassword=dio servername=brazil
serverpassword=brazil hladdress=193.1.1.11 l laddress=1500
```

This command will generate the a device configuration file for the Storage Agent, called `devconfig.txt`.

```
SET STANAME SA_DIOMEDE
SET STAPASSWORD 18dd2370
DEFINE SERVER BRAZIL HLADDRESS=193.1.1.11 LLADDRESS=1501 SERVERPA=185b801eb97c22
```

18.Finally, modify the option file that will be used by the TDP for MS Exchange client (`DSM.OPT`), so that it will use the Storage Agent rather than the Tivoli Storage Manager server (Brazil) to perform backups and restores. Add the following line to the option file:

```
ENABLELANFREE YES
```

19.Once this has been done, you can start the Storage Agent. It is advised to run this agent as a service. Refer to Section 5.3.6, “Installing the Storage Agent as a service” on page 139 for details.

Important

When using the Storage Agent, the filespaces on the server for the MS Exchange client will not change. This means that you can see MS Exchange backups created before the migration. However, if you try to restore these backups using the Storage Agent, the transaction will fail. In order to solve this problem, you should stop the Storage Agent service whenever restoring a pre-SAN backup. By doing so, you will reconnect to the Tivoli Storage Manager server using the LAN connection, and the restore will complete.

This completes our second example, in which we moved an entire server from a Windows NT Tivoli Storage Manager server, Diomede, to an AIX machine, Brazil. This was done using the export and import processes, porting all data and definitions from Diomede to Brazil. Finally, Diomede was re-configured so that it uses Brazil as Tivoli Storage Manager server for normal file backup and restore operations, TDP for MS Exchange, however, will move the MS Exchange backup data LAN-free to the shared library, using the Storage Agent.

Chapter 7. Serverless backup using TSM and SANergy

In this chapter we discuss SANergy concepts. We briefly describe SANergy setup, configuration, and usage with the Tivoli Storage Manager.

7.1 Tivoli SANergy File Sharing

In April 2000, Tivoli announced the product Tivoli SANergy File Sharing Version 2.1, which introduced LAN file sharing technologies to Storage Area Networks (SANs). In this section we describe the SANergy architecture and its cooperation with the Tivoli Storage Manager.

7.1.1 Tivoli SANergy File Sharing concepts

SANs are fast becoming the preferred method for companies to manage the vast amounts of data created by their e-business systems, creating networks to connect islands of information that can be shared quickly across the enterprise. SANs can link data centers and manage data across heterogeneous computing platforms for real-time backup and recovery, data migration, and resource sharing.

Tivoli SANergy File Sharing is the only SAN software that allows the sharing of application files and data between a variety of heterogeneous servers and workstations connected to a SAN. In addition, Tivoli SANergy File Sharing software uses only industry-standard file systems, enabling multiple computers simultaneous access to shared files through the SAN (compare Figure 102). This allows users to leverage existing technical resources instead of learning new tools or migrating data to a new file system infrastructure. This software allows SAN-connected computers to have the high-bandwidth disk connection of a SAN while keeping the security, maturity, and inherent file sharing abilities of a LAN.

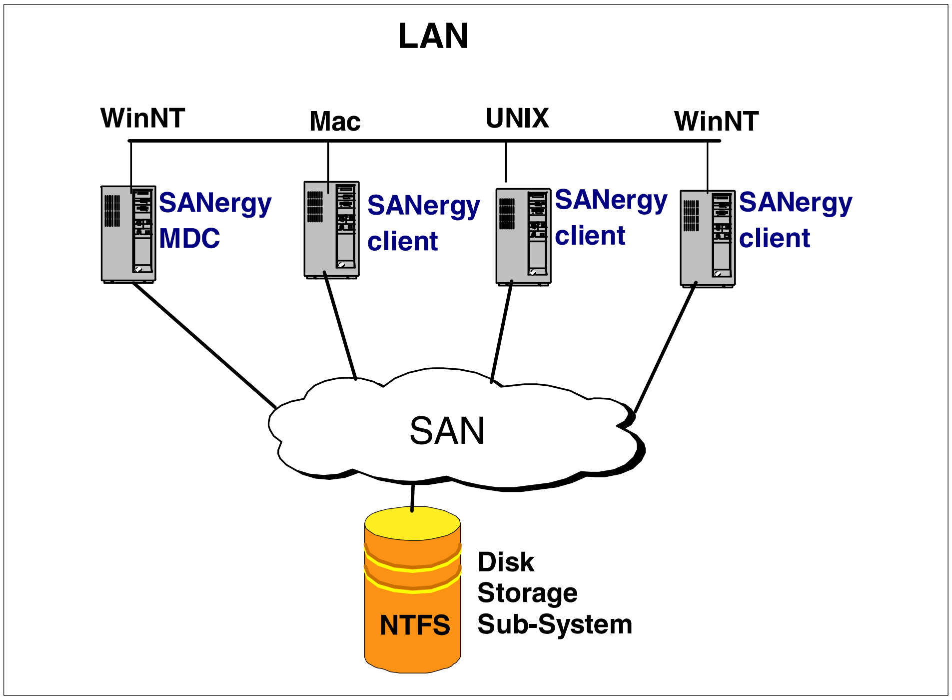

Figure 102. SANergy configuration

The SANergy software employs technology to combine the simplicity of LAN-based file sharing with the very high data transfer speeds afforded by modern Fibre Channel, SCSI, and SSA storage networks. This enables the use of high-speed, heterogeneous data sharing without the performance limiting bottlenecks of file servers and traditional networking protocols.

The Tivoli SANergyFS product is unique in that it extends standard file systems and network services provided by the operating systems that it supports (Windows NT, MacOS, and UNIX). As an O/S extension built on standard systems interfaces, SANergyFS fully supports the user interface, management, access control, and security features native to the host platforms, providing all the file system management, access control and security expected in a network. With SANergyFS, virtually any network-aware application can access any file at any time, and multiple systems can transparently share common data.

In addition to the SAN, Tivoli SANergyFS also uses a standard LAN for all the metadata associated with file transfers. Because Tivoli SANergyFS is NTFS-based, even if the SAN should fail, access to data via the LAN is still possible. Since each system has direct access to the Tivoli SAN-based storage, Tivoli SANergyFS can eliminate the file server as a single point of failure for mission-critical enterprise applications. Tivoli SANergyFS can also easily manage all data backup traffic over the storage network, while the users enjoy unimpeded LAN access to the existing file servers.

7.1.2 Tivoli SANergyFS architecture and data flow

The basic problem in storage area networking at the file level is keeping the separate operating systems up to date with each other's independent and asynchronous use of the storage. Tivoli SANergyFS is a hybrid of conventional networking and direct attached storage. Conventional networking is rich with abilities for keeping many computers coherent. That is, if one computer has an open view of a directory, and another changes that directory (adds/deletes a file), the view on all computers will change.

Conventional networking allows administrators to establish centralized access control lists and other data management facilities. Data "about" data is referred to as *metadata*. Examples include file names, file sizes, and access control lists. The Tivoli SANergyFS architecture lets metadata transactions take place over conventional LAN networking. The actual content of files moves on the high-speed direct SAN connection (Figure 103).

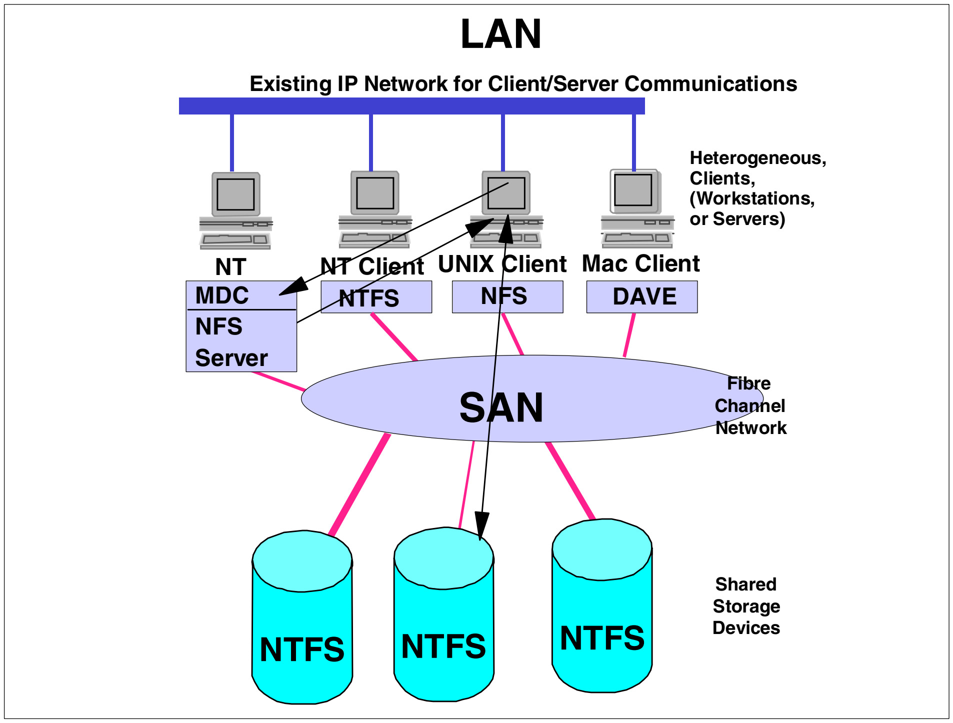

Figure 103. SANergy architecture and data flow

Tivoli SANergyFS works with Ethernet, ATM, or anything else that carries networking protocols. The network operating system can also be SMB (Windows NT), Appletalk, NFS (UNIX), or a combination. (Only certain protocols and combinations have been tested, however.)

Similarly, Tivoli SANergyFS supports any available disk-attached storage fabric. This includes Fibre Channel, SSA, SCSI, and any other disk-level connection. It is also possible for installations to use one set of physical wiring to carry both the LAN and storage traffic. (FC supports both IP and SCSI protocols.) When you use Tivoli SANergyFS, one computer in the workgroup is tagged as the *Meta Data Controller* (MDC) for a particular volume. You can have a single computer as the MDC for all volumes, or it can be spread around. The other computers are *SANergy clients* and use conventional networking to "mount" that volume, and Tivoli SANergyFS on those clients separates the metadata from the raw data automatically.

7.2 Configuring SANergy File Sharing

In Figure 104 we introduce our basic scenario for the Sanergy File Sharing project. We have set up SANergy between two hosts: one host running the Windows NT operating system, the other one running Windows 2000.

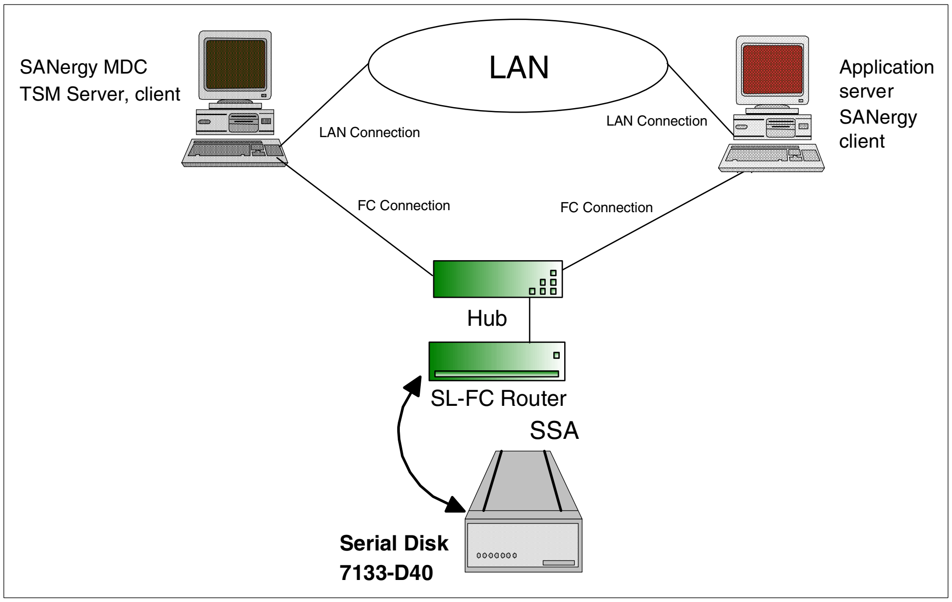

Figure 104. Preparing environment for SANergy FS scenario

We decided to configure the Windows NT machine as an MDC.

Note: If you are satisfied with connecting two FC hosts only, you also can connect these hosts directly to the SLIC router. In our environment, we planned to connect more hosts, and that is why we also used the SAN unmanaged hub.

Our next challenge was to test the Tivoli Storage Manager to back up SANergy volumes. For this purpose, we have extended our SAN environment as shown in Figure 105.

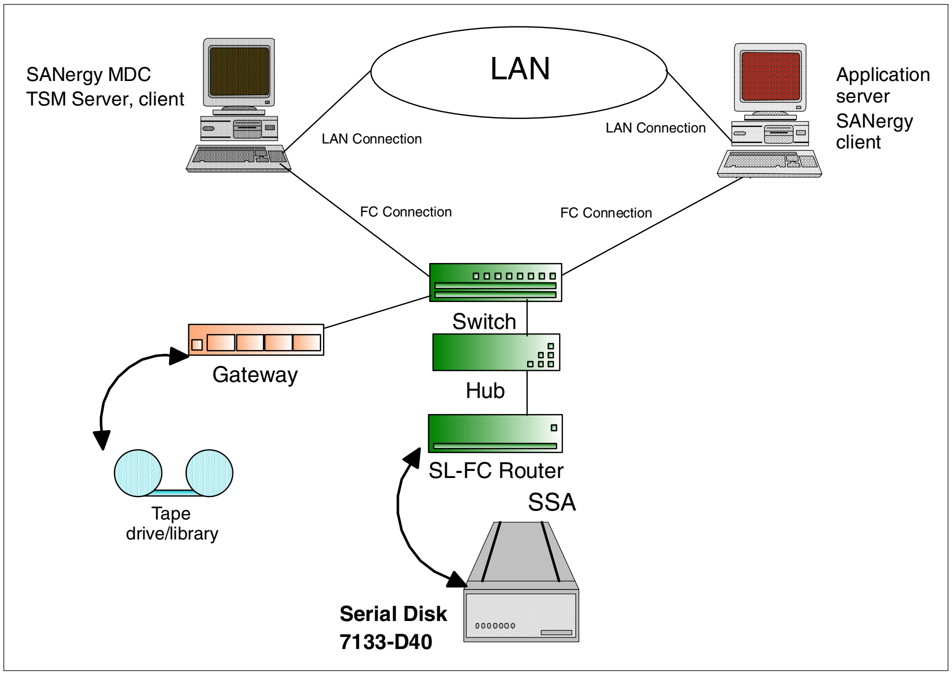

Figure 105. Preparing environment for TSM backup of SANergy volumes

We added the SAN Data Gateway 2108-G07 to our environment in order to connect the 3570 SCSI library to the SAN. For testing purposes only, we also used the 2103-S08 Switch.

7.2.1 Configuring SAN interconnectivity layer

There are basically three main components which are required to be configured:

- IBM Fibre Channel Switch 2109-S08

  For information on how to configure this device, read Section 3.2.4, “IBM Fibre Channel Switch 2109-S08 basic configuration” on page 69. Only basic configuration is required; there are no special considerations for SANergy.

- Fibre Channel Hub 2107-H07

  We describe this device in Section 3.2.1, “Fibre Channel Hub 2107-H07 basic configuration” on page 54. No special considerations are required to configure this device with the SANergy software.

- Vicom SLIC router

  Your SLIC router has to be able to communicate with the SSA disks and know their physical configuration. Please read Section 3.2.5, “Vicom SLIC Router FC-SL basic configuration” on page 72, where we describe all the necessary steps to configure this device. Then you can verify the configuration with the SLIC Manager software supplied with the SLIC router. We discuss the installation and configuration of the SLIC Manager software in Appendix B, “Configuring SLIC Manager on Windows” on page 273.

7.2.2 Verifying the host computer configuration

On your host computers, you need to configure the Host Bus Adapter to be able to connect this host to the SAN.

If you need more detailed information on how to configure HBAs in different operating system environments, please read Section 3.1, “Host Bus Adapters (HBAs)” on page 51.

Once your HBA is properly configured, you should be able to see the SSA disks which are connected to the SAN environment through the SLIC Router.

In Figure 106 you can see that the Disk Administrator sees the SSA volumes. You can also use the Disk Administrator to create partitions and format the volumes.

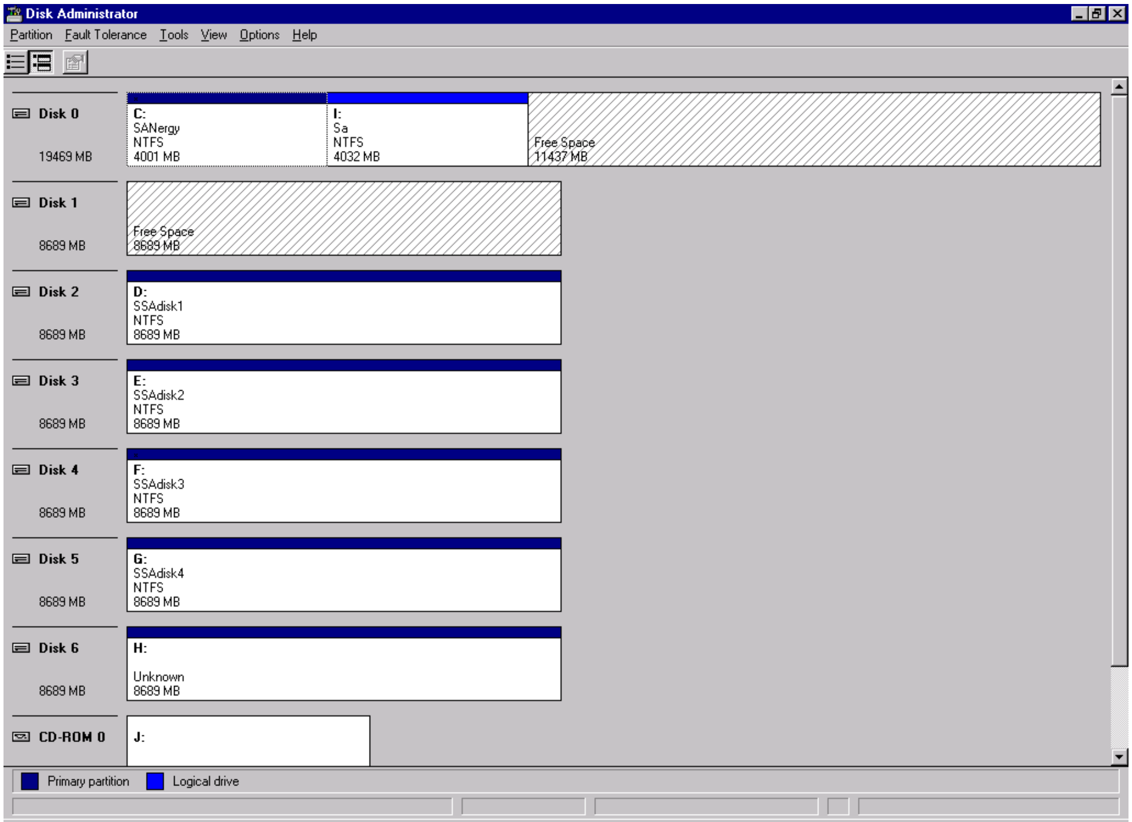

Figure 106. Preparing volumes using NT Disk Administrator

If you cannot see the SSA disks, verify your SAN environment — zoning and physical interconnection. Also try to check that the SLIC router is configured properly, but from our experience, once is the SLIC box configured, it works fine.

Note that the only way to check the SLIC router configuration is for you to install the SLIC Manager. See Appendix B, "Configuring SLIC Manager on Windows" on page 273 for detailed information on how to set up the SLIC manager software and how to use the basic commands to verify the SLIC router configuration.

7.2.3 Prepare your volumes with the NT Disk Administrator

Run the Disk Administrator utility on the Windows NT host you plan to configure as the MDC and re partition and format the SSA disks. Note that under Windows NT it is only NTFS file system supported.

7.2.4 Install and configure SANergy software

In this section, we give you a brief overview on how to set up the SANergy MDC machine, and what steps are necessary to go through on the SANergy client side.

For more information about the SANergy configuration, read the documentation supplied with the product and/or the IBM Redbook, *Tivoli Storage Manager 3.7.3 & 4.1 Technical Guide,* SG24-6110

7.2.4.1 Configuring SANergy MDC

After you have installed the SANergy software from the installation media, a SANergy configuration tool appears. Now you have to assign volumes to the MDC (see Figure 107).

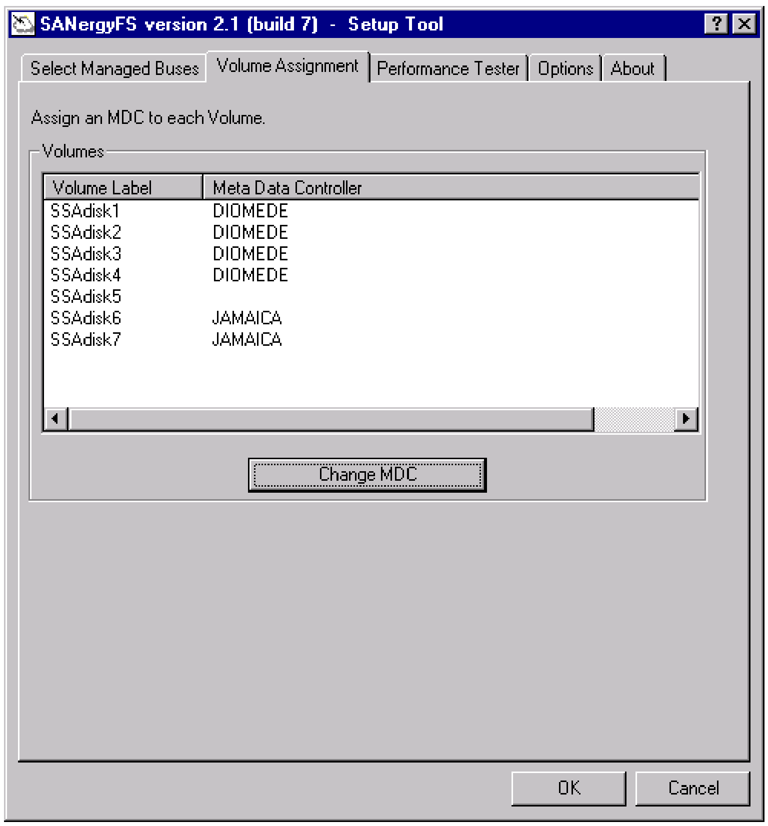

Figure 107. Assigning volumes to the SANergy MDC

As shown in Figure 107 on page 229, click on the Volume Assignment and assign all volumes you want to belong to your MDC. In our scenario, we had two MDC machines — Jamaica and Diomede. Also note that the SSAdisk5 does not belong to any MDC machine.

You may now want to verify if you are able to read/write the data to your SANergy volumes. If you click on Performance Tester in the SANergy configuration tool, you can perform the read/write operation.

Define NT drive sharing for a SANergy volume

Now you need to define the NT share for your SANergy volumes on the MDC. This gives your other hosts — SANergy clients — the possibility to “mount” the SANergy volumes. If you also have UNIX clients in your environment, you have to install an NFS server on the MDC host. UNIX clients can only use NFS to mount the SANergy volumes. Figure 108 illustrates how to define the sharing for a SANergy volume.

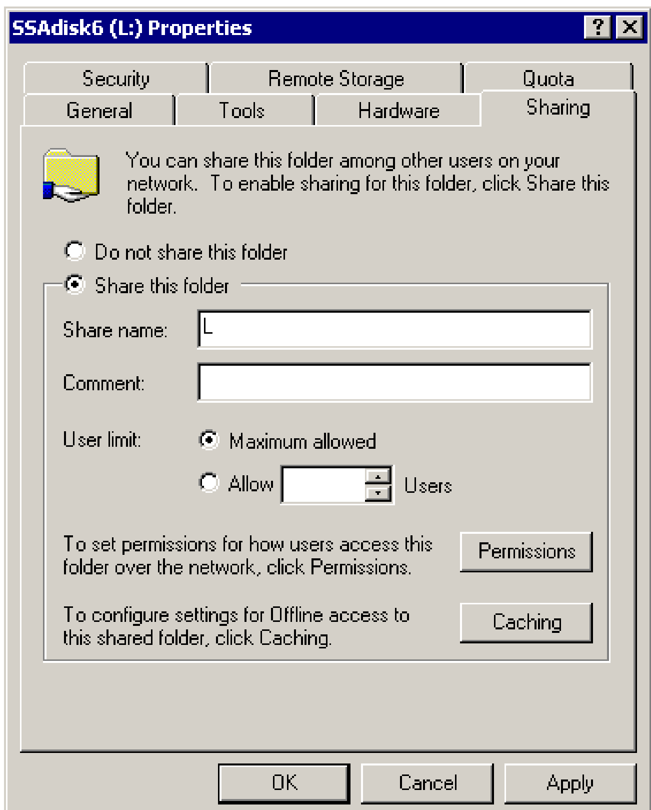

Figure 108. Define NT share for a SANergy volume

7.2.4.2 Configuring SANergy clients

Now install the SANergy software on your client computers. In SANergy Version 2.1 there is no separate code for SANergy MDC and SANergy clients. After the installation is done, reboot the computer if it is required by the operating system.

If your operating system environment is Windows NT and your SANergy volumes are part of a stripe group, you have to restore the configuration from the disk where you have saved this configuration using the NT Disk Administrator.

Now mount the MDC volume using regular Windows networking tool. If you are configuring the SANergy client in the UNIX environment, mount the volume using the `mount` command.

It is recommended to use the following options for NFS mounting in the UNIX environment:

```
mount -o acregmin=0,acregmax=0,actimeo=0,nfs,rw host:/shared/volume /mnt
```

7.3 Using TSM with SANergy

Using SANergy together with the Tivoli Storage Manager will give you the possibility to transfer your data through the SAN. It supports both LAN-free and serverless types of backup/restore, in both cases the data transfer will be off-loaded to the SAN.

In the following sections, we discuss two possible configurations.

7.3.1 TSM backup/restore using SANergy — scenario1

In this example, the application hosts are running SANergy client code, are mounting the disks they need, and are sharing data among themselves. On our SANergy MDC machine we also have the Tivoli Storage Manager server running. This MDC machine owns all of the disk volumes and these volumes can therefore be accessed locally, so we are able to back up the data on these volumes with the Tivoli Storage Manager backup/archive client. Restores are also performed locally.

In Figure 109 you can see our SANergy/Tivoli Storage Manager scenario.

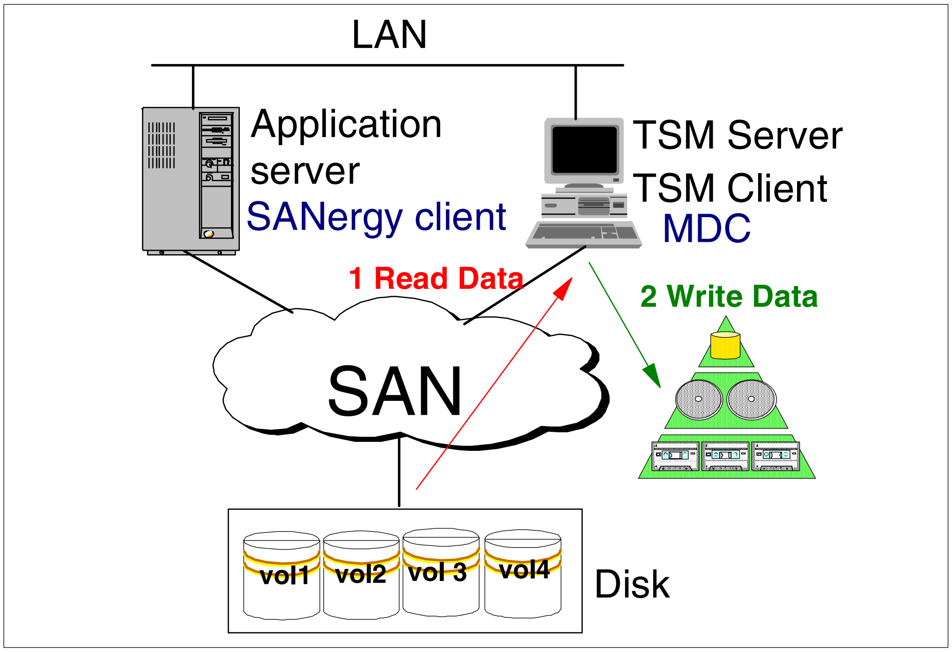

Figure 109. TSM using SANergy — scenario 1

Note that in this case, no data is transferred through the LAN, not even metadata. That is because there is no backup/restore action on your application server.

7.3.2 TSM backup/restore using SANergy — scenario2

In our second scenario, the Tivoli Storage Manager server is installed on the application system, which is a SANergy client. This client mounts its volumes from the MDC machine. Figure 110 shows the scenario in our lab.

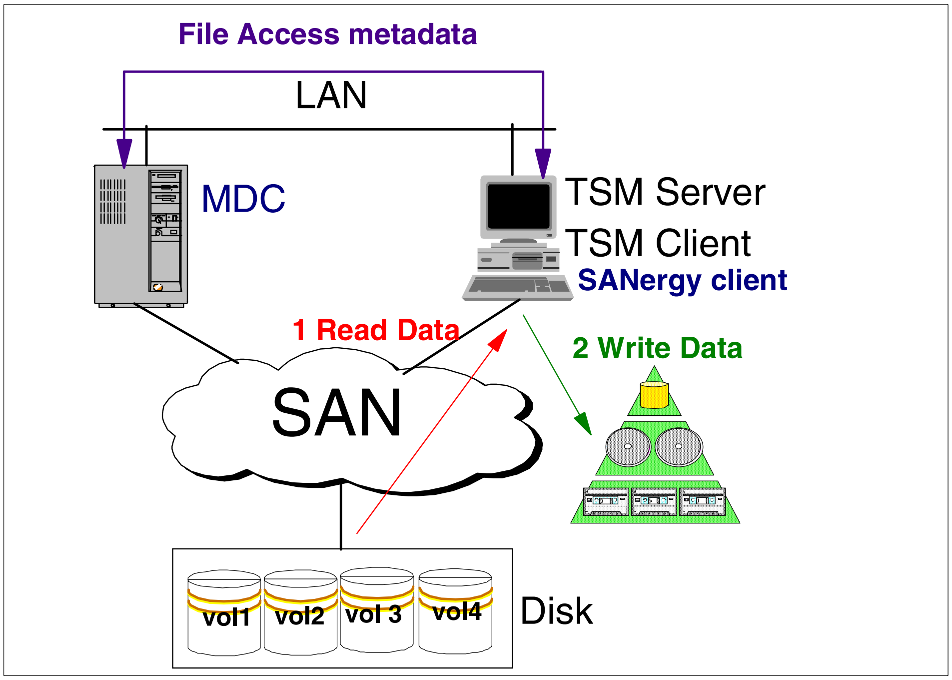

Figure 110. TSM using SANergy — scenario 2

When the Tivoli Storage Manager client begins to back up the data, it will need to get the metadata from the MDC machine. For this purpose, the TCP/IP transport over the LAN will be used. But the raw data will be still transmitted through the SAN.

Note that this scenario requires a homogeneous environment to do successful restores with all file metadata, which means that the platform type of the system running the Tivoli Storage Manager client has to be compatible with the MDC systems.

7.4 SANergy usage considerations

If you have more than one MDC configured, and need to assign a volume to another MDC, you have to remove the volume from the current MDC member list first. Then you can assign this volume to the other MDC. That is because the SANergy software currently does not let you change the membership of an MDC volume directly.

If you add a new volume to the MDC and define the sharing for this volume which allows this volume to be “mounted” on the client system, then you can just mount this volume on the client system. It is not necessary to restart the SANergy MDC service and/or reboot the client computer system. You can also delete a volume from the MDC without the need to restart SANergy.

In SANergy Version 2.1, when you run an MDC on the Solaris machine, this MDC can only be used by UNIX SANergy clients. The Windows platform is not supported. That is because UNIX does not support the Windows-like volume sharing, so you are not able to map a drive from a UNIX machine. The possible work-around would be to install Samba-like software on UNIX. But keep in mind that this may not work, and that it is not supported.

On the other hand, your Windows NT or Windows 2000 MDC can serve the UNIX clients, but this requires you to install an NFS server on the Windows system. That is because UNIX will use NFS to mount SANergy volumes from your MDC machine.

In our environment, we have tested Hummingbird’s NFS Maestro Server and Esker’s Tun Plus NFS server. For more information about these products, see following Web pages:

- `www.hummingbird.com/products/evals/index.html`
- `www.esker.com/Q300/products/host_access/tun_plus/2.1.2_tunplus.html`

When running MDC and Tivoli Storage Manager server on one machine, there may be a higher demand on the CPU performance. Consider the number of clients you want to connect to your MDC and the amount of data you plan to back up, and make sure your hardware is efficient enough.

7.5 Considerations on using TSM in a SANergy environment

Consider if you need to back up open files to the Tivoli Storage Manager serve; this means if you need to have set the SERIALIZATION parameter in your copy group definition to dynamic or shared dynamic. In many cases, the data backed up in such a way may be inconsistent. This applies especially to the database files.

Do not use SANergy for backing up online databases. For this purpose, use Tivoli Data Protection products together with Storage Agent software, which gives you the possibility to back up your databases online through the SAN. We are aware that today there are only two TDP applications (TDP for MS Exchange and TDP for R/3) supported, but it is intended to support the other TDP products soon.

Chapter 8. Availability considerations in a SAN

This chapter discusses solutions and issues related to the usage of the SAN fabric as a tool to improve the availability of the Tivoli Storage Manager installation. As introduced in Section 2.2.1, “Availability” on page 38, SAN implementations are often viewed as high availability solutions. In this chapter, we will consider these proposed solutions, and discuss the way they can be used when combined with Tivoli Storage Manager and tape libraries. These solutions include:

- Remote tape attachment
- Multiple paths to a tape device

We will look more closely at these solutions, documenting their use and possibilities within a Tivoli Storage Manager SAN environment.

8.1 Remote tape implementations

As stated in Section 2.2.1.1, “Remote tape library connection” on page 38, one of the advantages of Fibre Channel attached tape storage is the increased distance between the library and the Tivoli Storage Manager server.

Whereas normal SCSI implementation have a limit of 25 meters (using differential SCSI), Fibre Channel connections can reach 10 kilometers. Figure 111 shows such a setup. Using a switch and a hub, both equipped with Long Wave GBICs, a Magstar 3590 E11 attached autoloader is attached to the Windows NT host.

The Magstar 3590 model E11 is a rack-mounted tape driver with an Automatic Cartridge Facility (ACF). The 10 cartridge ACF allows the drives to be used as a mini-library with access to up to 1.2 TB of data. The drive is directly connected to the SAN using the Fibre Channel attachment feature. To connect the tape and the HBA to the fabric, short wave cables are used. The connection between the switch and the hub is done using long wave cables.

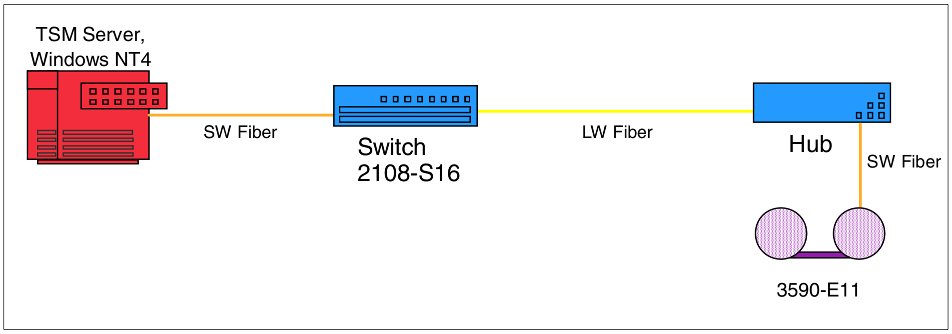

Figure 111. Distant connection of a 3590 FC attached tape device

Various implementations that can be done using this setup, for example, automatic vaulting. By placing the library in an off-site, secured location, the need for manual export of tape volumes disappears. However, the need to set this up as a copy storage pool of a local primary library does not. Copy storage pools have a second use besides disaster recovery. This use, and probably the most important one, is recovering from tape volume failures.

The next example shows the setup of this library. To define this tape device to the Windows NT server, use the normal tape and library definitions. Although the library is directly FC connected, it is seen as a SCSI device. Figure 112 shows that the 3590-E11 is seen as library device lb0.1.0.3 and tape drive mt0.0.0.3.

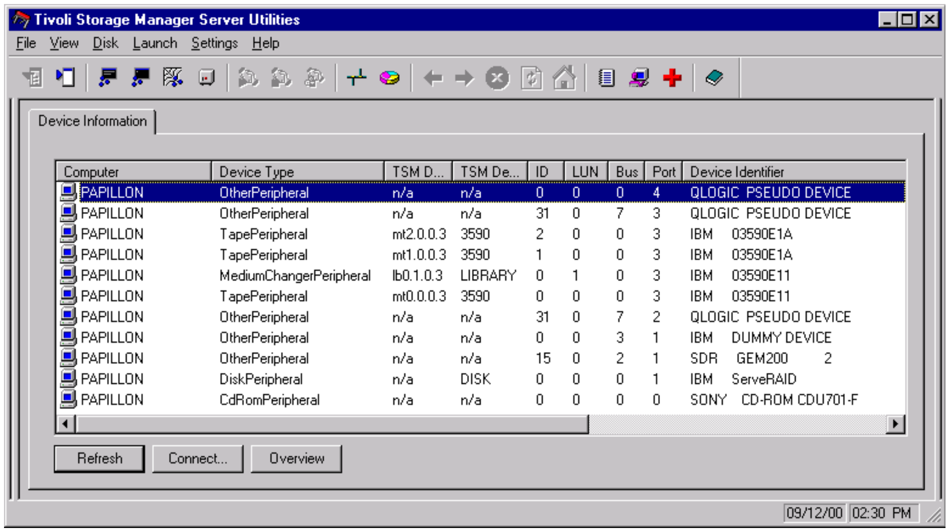

Figure 112. SCSI devices seen by the ADSMSCI driver

To define these devices, issue the following commands:

```
DEFINE LIBRARY 3590 libtype=scsi device=lb0.1.0.3
DEFINE DRIVE 3590 3590drive device=mt0.0.0.3 element=16
```

After defining these devices, you will still need to define a device class, and storage pools using this device class.

8.2 Fabric failover

One of the features of a SAN fabric is its potential to deal with component failure. The main reason for this is the possibility to create multiple paths between the devices and the hosts. Whereas SCSI implementations have typically one or two connections between the host and the device, a switched fabric can almost create an unlimited number of connections between the host and the device. This is done by combining switches, using *Inter-Switch Links* or *ISLs.* An example of such a topology is shown in Figure 113.

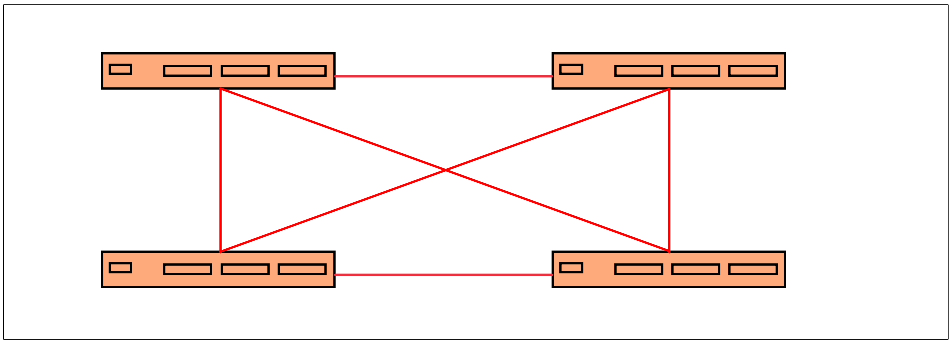

Figure 113. Meshed switch topology

A fully meshed topology as shown above has two advantages:

- Minimized latency due to the fact that every switch is only one hop away from the other one. This increases performance.
- The impact of link and switch failure is minimized due to path and switch duplication.

In SAN publications, it is often stated that a multi-switch fabric offers more flexibility to build a fault tolerant system. The path selection protocol allows multiple links to operate at the same time between any two switches in the fabric, so multiple redundant paths can be created. The extra links do not have to be reserved and maintained idle: all links carry traffic. If a link goes down, the traffic that was carried over that link will be simply transferred to the other link. This transfer takes places automatically, with no human intervention. Even if a switch goes down, all other switches in the fabric, and the end nodes connected to them, are not affected.

Whereas this certainly true from a fabric point of view, it is important to know how an actual user of the SAN fabric will react. In our case, this user can be identified as a Tivoli Storage Manager server or Storage Agent transferring its data to a SAN attached SCSI tape device.

Note

Direct Fibre Channel attached devices, like the 3590-E11 drives and loaders, are recognized as SCSI devices. Therefore, the discussion here applies to these types of devices as well.

The following sections examine the specifics for a Tivoli Storage Manager server using a SCSI devices attached to the SAN. Two main problems exist:

- SCSI ID duplication
- SCSI bus reserves

We will look at these two problems in more detail, and propose solutions for them.

8.2.1 SCSI ID duplication

SCSI devices, attached through the SAN using the San Data Gateway are used by the application using their virtual SCSI address. This virtual address is actually dynamically created by the SAN fabric, and has no relation with the actual device SCSI address. Each path to the device will be proposed as a SCSI connection to a device.

8.2.1.1 The problem

The result of these virtual addresses, as shown in Figure 114, is that when there are multiple paths to one device, the ADSMSCSI.SYS device driver sees these paths as connections to different devices. This means that if you have two paths to a tape drive, you will see two tape drives, each with their own SCSI address.

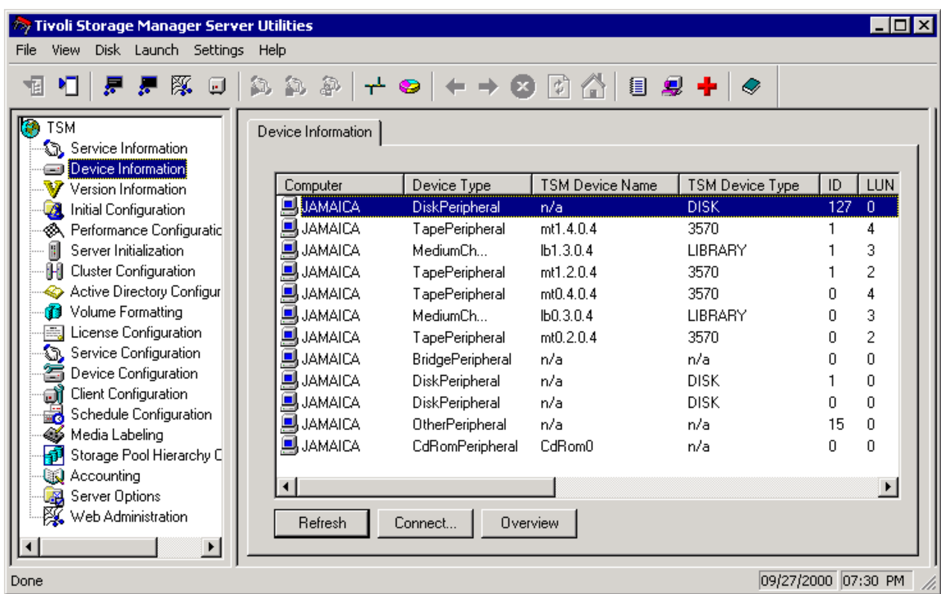

Figure 114. SCSI devices seen in a switched fabric

Whereas this is not a problem initially (you could define only one drive using the SCSI ID you need), the problems appear when the machine or fabric is rebooted. As stated above, SCSI addresses are allocated dynamically when the system connects to the SAN. Under normal conditions, this means that you will always have the same SCSI address for a certain path. However, under abnormal conditions (for example, when one path is not available), the SCSI addresses will change. This means that you have to reconfigure all drive and tape definitions on your Tivoli Storage Manager servers.

8.2.1.2 The solution

The solution to this problem is *zoning*. Using zoning, you can enable certain connections and disable others. This means that instead of having two paths available at the same time, you will have only one. The result is that you will see only one SCSI device (which is the reality). When a connection breaks, the switch configuration can be changed, so that it activates the secondary path. The SCSI ID will remain the same after the configuration switch, when only one HBA adapter is present. So, you do not need to reconfigure your devices at a Tivoli Storage Manager level.

An example of how to setup zoning, and use it in case of a failover, is given in Section 8.2.3, “Cascading switches implementation” on page 241.

8.2.2 SCSI bus reserve

A second, more severe problem, when dealing with SCSI attached devices, involves the SCSI bus reserve.

8.2.2.1 The problem

Whenever a connection is established between a Tivoli Storage Manager server and a SCSI tape drive or the library robotics, the SCSI bus is reserved for that connection. This means that only that connection can use the SCSI device, and no other hosts can connect to the device. More important, however, is that only the host that opened the connection can close it. This implies that whenever a Tivoli Storage Manager server uses a tape drive, it and only it can reset the device, and make it available for other servers.

When for some reason the connection is stopped, and cannot be reestablished using the exact same path and server, the device stays in use. The only way to recover is to perform a hardware reset of the SCSI bus. In the case of a SAN implementation, this means at least resetting the San Data Gateway. However, since this device is used by all library clients and the library manager, resetting it will disrupt all SCSI connections.

8.2.2.2 The solution

In the following topics we will look at two situations, and document the behavior in case of a fabric failure. We will begin with setting up zoning to block the multiple SCSI addresses, and document the steps required to keep working with the Tivoli Storage Manager servers after this failure. Since this behavior is dependent on the type of failure and the moment it occurs (defining the state a device is in on a server), the examples and recovery steps shown below are not exclusive.

The standard rule that can be adopted whenever dealing with SAN failures or problems, is to reset the entire implementation. This means resetting the SCSI devices, the San Data Gateway, the hubs and switches, and finally the Tivoli Storage Manager servers. The order of this reset should be as indicated above.

8.2.3 Cascading switches implementation

The first setup will be done using two switches, linked together using an inter-switch link (ISL). Every host has only one connection with one of both switches. Whereas in this case, the error recovery is only available for link failure or switch port failure (only the port to the SAN Data Gateway), it is a good base example to document the reaction of the components after failure.

In case of failure of the switch port connected to the SAN Data Gateway, the connection can be reestablished using the ISL, resulting in the two switches cascading. The first important limitation of this setup is that currently, AIX does not support switch cascading. This means that we will only look at a failure on the switch connected to the library clients (Windows NT and Windows 2000). Figure 115 shows this configuration.

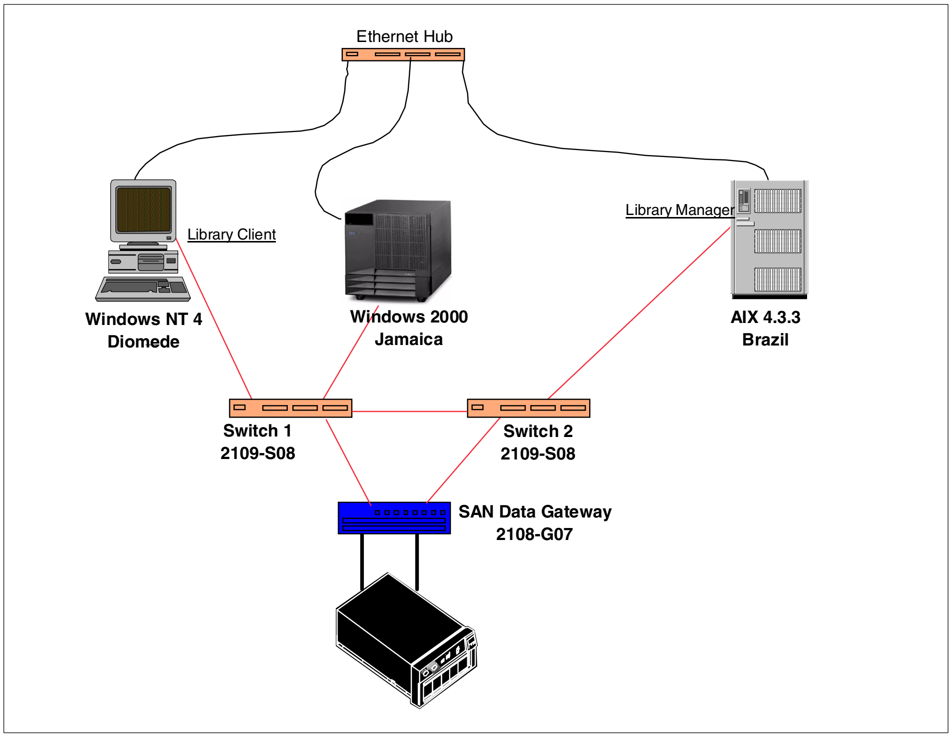

Figure 115. Cascading switches setup

8.2.3.1 Setting up the zoning

Before using the above setup, we must define a zoning scheme for our switches, that blocks multiple concurrent paths to the SCSI device. If not, we would see the same devices multiple times (see Figure 114 on page 239). And, since SCSI address allocation is dynamic, we can never be sure which SCSI address relates to which path.

The first concept in the zoning configuration involves the zone aliases. A zone alias can be used to give a switch port or group of ports a certain name. Although this is optional, it is easier to work with and manage than using the actual port names for setting up the zones. In our setup, we decided to create aliases based on the machine that is connecting to a certain port. Using the Storwatch Fibre Channel Switch Specialist software, we started the Zone Administration component.

Figure 116 shows the definition of an alias `Jamaica`, using port 1 on `switch1`, belonging to the first switch domain. The naming convention for these port assignments is (domain number, port number). As an alternative, you could define aliases using the *world wide port name* (WWN) of the switch port that you are using.

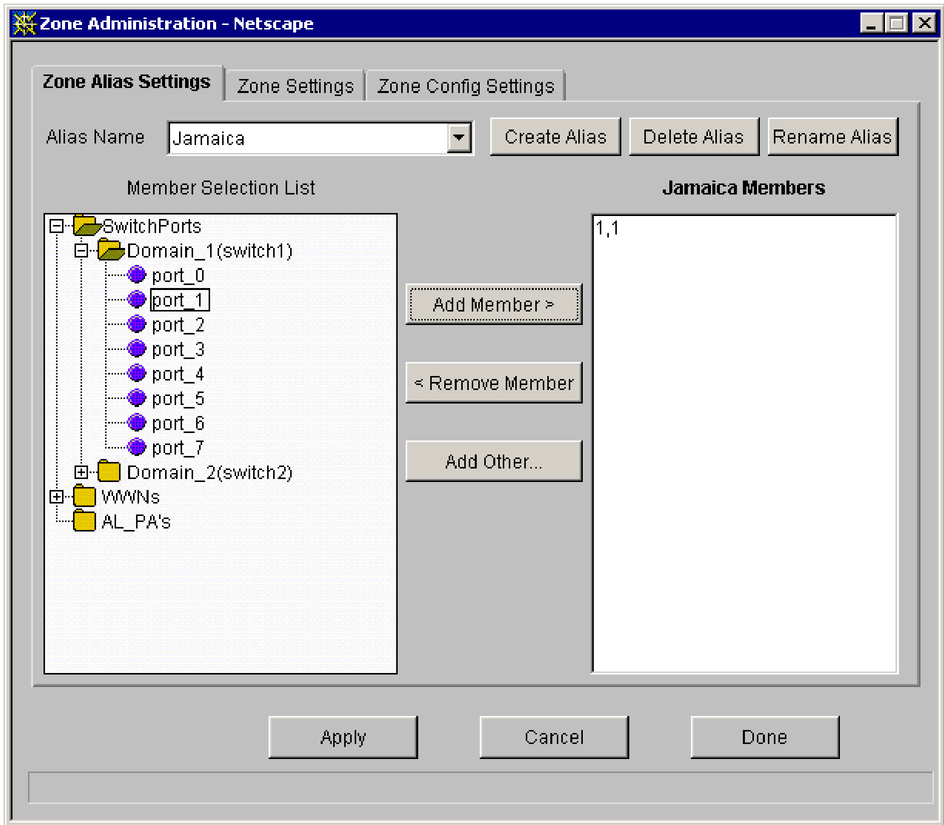

Figure 116. Setting zone aliases

These steps were repeated for all connections that were in use for our current test. A special alias was the `ISL` alias (see Figure 117). Since the ISL always uses two ports in our switch configuration, we defined it as such, meaning that the ISL alias related to two ports instead of one.

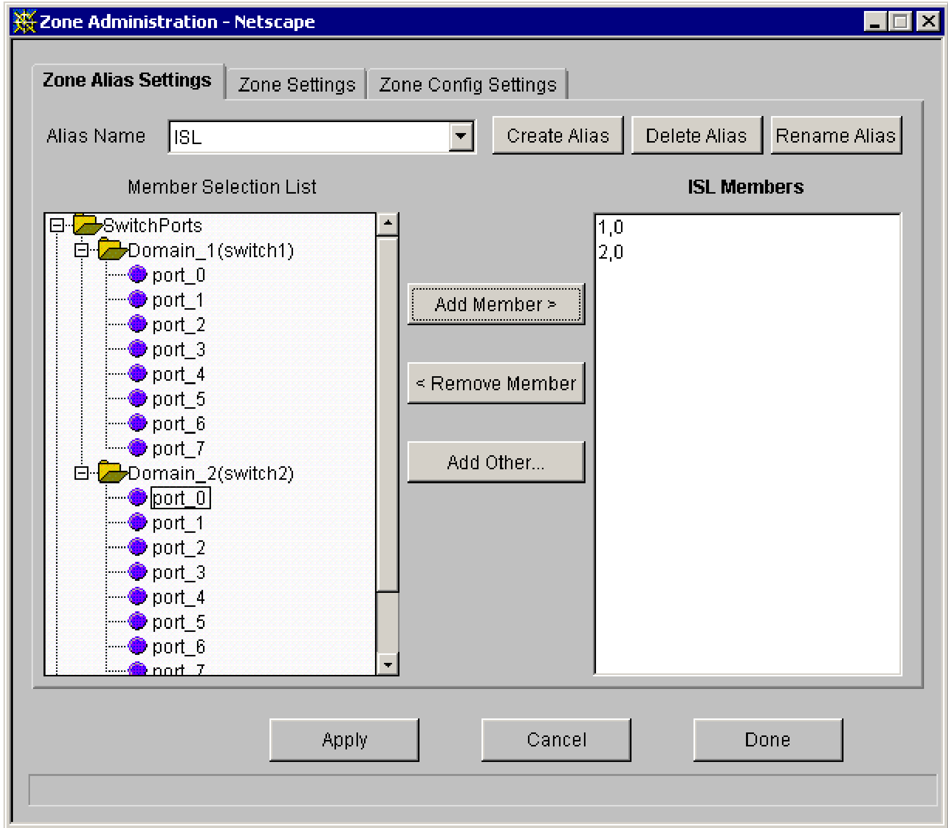

Figure 117. ISL alias definition

Next, it is time to define the zones and their members. Again, using the zone administration part of the Storwatch software, we defined the required zones. The logic behind these definitions is as follows: the purpose of defining these zones is to create virtual links between ports, so that at any moment, only one connection exists between the Tivoli Storage Manager servers and the tape devices. In addition, alternate paths must be available when the normal path is unavailable.

Following this rule, we defined three zones:

- A zone connecting Jamaica and Diomede through switch one to the San Data Gateway
- A zone connecting Brazil through switch two to the San Data Gateway
- A zone connecting Jamaica and Diomede through switch one and two to the San Data Gateway, using the ISL.

Figure 118 shows the actual definition of a zone. In Zone_0_1, Jamaica, Diomede, and the first SAN Data Gateway channel are grouped.

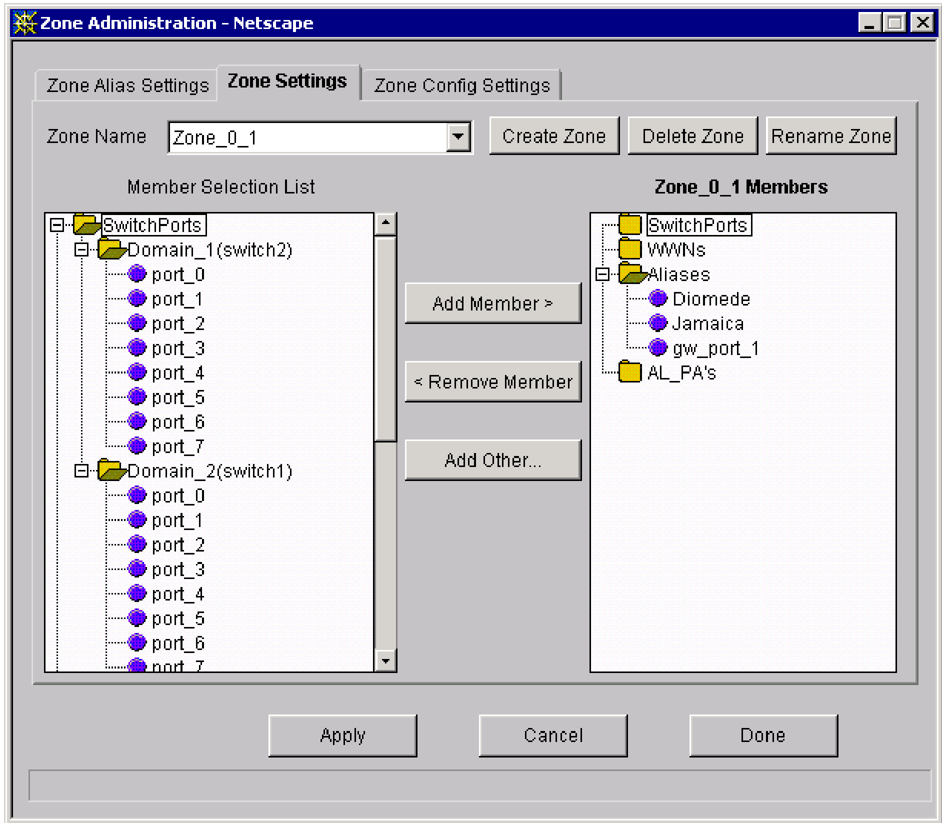

Figure 118. Defining a zone

Three zones are created, `Zone_0_1`, `Zone_0_2`, and `Zone_1_1`. Figure 119 shows which ports are members of which zone.

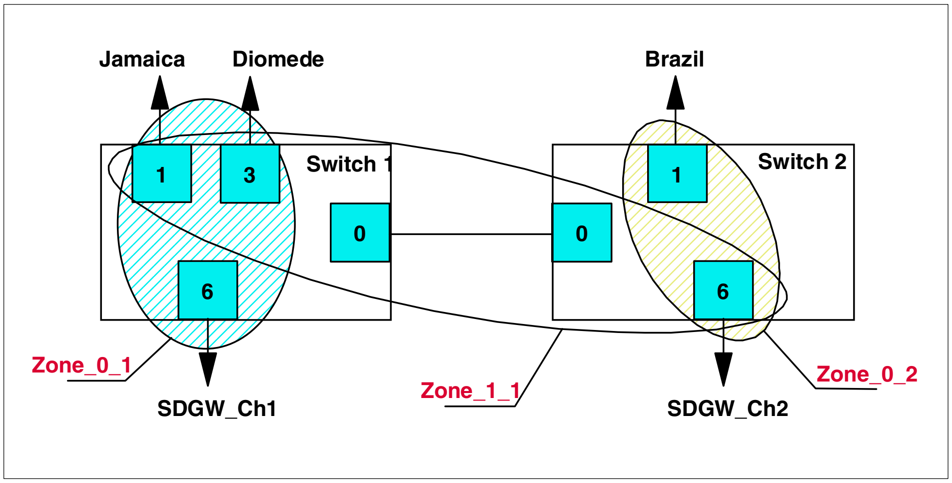

Figure 119. Zone definitions

After defining the zones, the final step is activating the zone configuration. Zone configurations are constructed using groups of zones. More than one configuration can be defined, but only one can be active. For our situation, two configurations were defined:

Normal This configuration is used in normal conditions. `Zone_0_1` and `Zone_0_2` are member of this configuration. This configuration is the active one.

Failover The second configuration will be used in case of problems, and uses `Zone_1_1` and `Zone_0_2`.

After defining and activating the configuration, the zoning is complete. To check the zoning, you can open a telnet session to the switch, and issue the `CfgShow` command.

```
switch1:admin> cfgshow
Defined configuration:
 cfg:   Normal Zone_0_1; Zone_0_2
 cfg: Failover Zone_1_1; Zone_0_2

zone: Zone_0_1
              Diomede; Jamaica; San_Data_Gateway_FC_1
zone: Zone_1_1
              Diomede; Jamaica; ISL; San_Data_Gateway_FC_2
zone: Zone_0_2
              Brazil; San_Data_Gateway_FC_2

alias: Brazil            2,1
alias: Diomede           1,3
alias: ISL               1,0; 2,0
alias: Jamaica           1,1
alias: San_Data_Gateway_FC_11,6
alias: San_Data_Gateway_FC_2 2,6

Effective configuration:
 cfg:   Normal
zone:  Zone_0_1
          1,3
          1,1
          1,6
zone: Zone_0_2
          2,6
          2,1
```

Note

When setting up zoning, you should be sure that your definitions are distributed to all switches. The inter-switch communication is provided by the ISL. To check that you have access with all switches, they should appear on the start windows of the Storwatch application. If they do not, there is a problem. If switches are added afterwards, the zone configuration will be distributed to the new switches, unless conflicts exist. For more detailed information on zoning, please refer to the IBM SAN Fibre Channel Switch 2109-S08 and S16 product documentation.

The result of the zoning for our Tivoli Storage Manager implementation is shown in Figure 120.

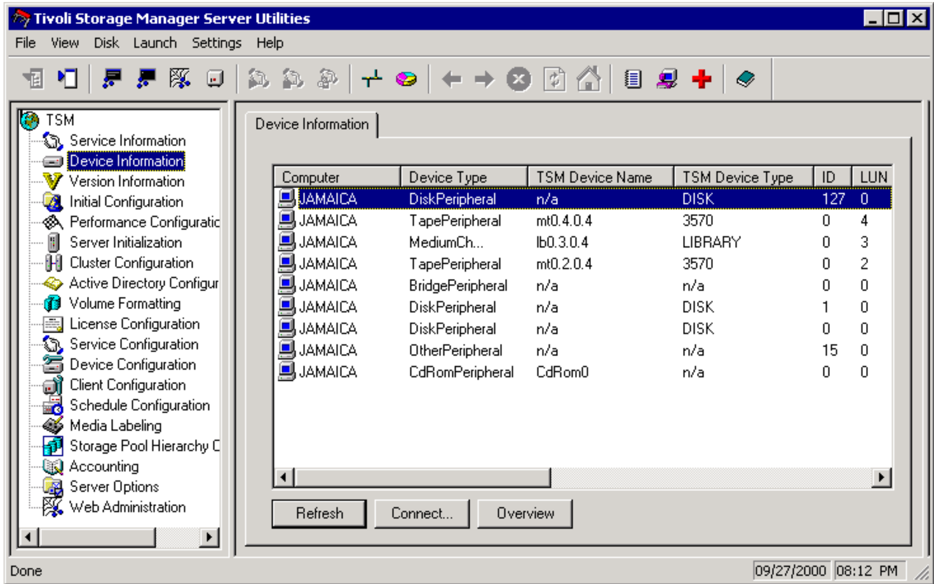

Figure 120. SCSI devices seen after zoning

8.2.3.2 Recovering from a link failure

The next step in our process will be the behavior of the components when a failure occurs. The failure we simulated was a broken link between switch 1 and the San Data Gateway. This was done by disabling port 6 on switch 1. In our test case, we used Jamaica as library client, and Brazil as library server (see Figure 121).

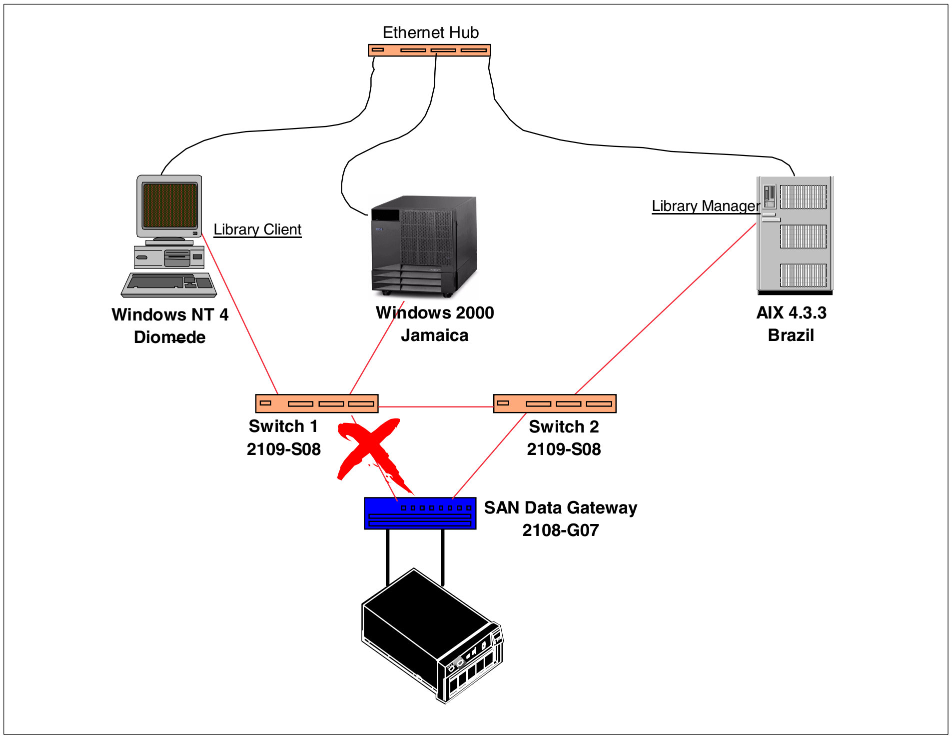

Figure 121. Link failure test setup

Failure indications

The first step in recovering from a failure is knowing that the failure occurred. The errors encountered on Jamaica after provoking the failure were not consistent. They ranged from server reboots, time-outs on drives to no error messages at all. One pattern that could be detected was that sessions or processes that are in progress do not generate errors. For example, a backup session that is using the tape drive when it fails, will remain in a running state. No data is actually moved, and, due to the fact that the state of the session is running, it will not time-out. Processes or sessions that request tape volumes after the error was generated will generate time-out errors.

The above leads us to conclude that when implementing a SAN, management of the SAN is very important. One feature of the switch is that it can generate SNMP traps, which can be forwarded to a systems management environment.

The situation after the failure

After trapping the error condition, it is now time to try to recover the SAN with the least impact on other systems. The first step is to determine what state the different components in our library sharing environment are in. The following error conditions were encountered:

- When the library client loses its SAN connection to the tape device, the library manager is not aware of that. This means that every mount request send by the library client to the library manager will result in a tape drive that is mounted.
- The Tivoli Storage Manager server that had a connection to the drive when the SAN component failed, is unable to release the drive. Due to this SCSI bus reserve, we are unable to access that drive using any other connection.
- Finally, there is the problem of the inconsistent state the Tivoli Storage Manager Library Manager has for the tape device. The library manager will see this device as free or loaded, but without knowing to whom it is allocated. As a result, any operation done to this device will result in an I/O error.

Steps to recovery

The first step when recovering is trying to stop all tape requesting processes and sessions on the machine that was connected through the failing link. Also, you need to ensure that no new ones can occur. This can be done by disabling the library client for client sessions, and deactivating all administrative schedules that might run in the near future.

Next, try to free the drives that the library client has in use. This can be done by canceling the processes and sessions that are reserving the drive. Since the dismount is done by the library manager, this can be successful. However, realize that the processes are in an error condition, and might take a while to stop.

However, just in case freeing the drives fails (by reboot of the library client, for example), the drive or drives that were used will be in an open state, and cannot be reset by the library manager. To check if this situation has occurred, you can use the `SHOW LIBRARY` command on the server. Below, the output of this command is shown. In our case, Drive 0 has a loaded state, but is not allocated to a certain Tivoli Storage Manager server. Also, the drive is physically loaded by a tape volume.

Note

The show library command is given for problem identification purposes only. Since this command is not supported for normal use, you should use it only when directed by Tivoli.

```
Library 3570LIB (type SCSI):
  refs=0, online=1, borrows=0, create=0, update=0
   driveListBusy=0
  libvol_lock_id=0, libvolLock_count=0, SeqnBase=0
  Library is shared.
  library extension at 3138AE50
  autochanger list:
    dev=/dev/rmt2.smc, busy=0, online=1, LTS=0

  drive list:
    DRIVE0 (/dev/rmt2):
      Device Class = 3570,
      RD Capabilities = 000C0000,
      (3570C,3570B),
      WR Capabilities = 000C0000,
      (3570C,3570B),
      Drive state: LOADED, Drive Descriptor state = NO PENDING CHANGES,
      online = 0, polled = 0,
      allocated to (), volume owner is (0863D3),
      priorReadFailed = 0, priorWriteFailed = 0, priorLocateFailed = 0,
      update count = 0, checking = 0,
      offline time/date: 13:05:46  2000/09/30
      Clean Freq(GB)=-1, Bytes Proc(KB)=0, needsClean=0.
      Log sense counters last dismount:
    SCSI-specific fields:
      element=16
      isTwoSided=0, whichside=0
      loaded volume home slot=45
      loaded volume name=0863D3

    DRIVE1 (/dev/rmt7):
      Device Class = 3570,
      RD Capabilities = 000C0000,
      (3570C,3570B),
      WR Capabilities = 000C0000,
      (3570C,3570B),
      Drive state: LOADED, Drive Descriptor state = NO PENDING CHANGES,
      online = 1, polled = 0,
      allocated to (DIOMEDE), volume owner is (083573),
      priorReadFailed = 0, priorWriteFailed = 0, priorLocateFailed = 0,
      update count = 0, checking = 0,
      offline time/date: 00:00:00  1900/00/00
      Clean Freq(GB)=-1, Bytes Proc(KB)=0, needsClean=0.
      Log sense counters last dismount:
    SCSI-specific fields:
      element=17
      isTwoSided=0, whichside=0
      loaded volume home slot=47
      loaded volume name=083573
```

As a result of this situation, all attempts to use the drive will result in an I/O error. The solution is to place the failing drive in an offline state on all Tivoli Storage Manager servers sharing this library, using the `UPDATE DRIVE` command with the ONLINE=NO parameter. This will prevent attempts to use the drive, and therefore will stop the I/O errors. If this is not done, Tivoli Storage Manager will set all volumes on, for which it receives an I/O error when trying to mount them in a read/write state.

After placing the Tivoli Storage Manager servers in a state in which they can continue working, it is now time to reinsert the Tivoli Storage Manager server that used the failed connection. First, update the switch configuration, activating the configuration that you created for failover situations. Next, reboot the Tivoli Storage Manager server. This is necessary to initialize the SCSI connection. After the reboot, the server should be able to connect the library again.

Keep in mind that the library is in an inconsistent state. Further recovery is required, but can be postponed until no other servers need to use the library. This recovery will include resetting the Tivoli Storage Manager servers, and the SCSI bus.

Start by stopping all Tivoli Storage Manager servers, ensuring that they do not use the library at that moment. After they have halted, remove any tape volumes mounted in the drives and reboot the San Data Gateway. This will reset the SCSI bus. After the SCSI bus has been reset, you can restart the Tivoli Storage Manager servers, and resume normal operations using the switch backup configuration.

In case the switch failure occurs when the drives are not in use, recovery is much easier. You only need to reconfigure the switch, and reboot the Tivoli Storage Manager servers for which new configuration changed the path to the devices.

8.2.4 Dual Host Bus Adapters

A second situation we will look at is an implementation in which each system has two HBAs connected to each switch (see Figure 122). This implementation could help in case of host bus adapter or switch failure.

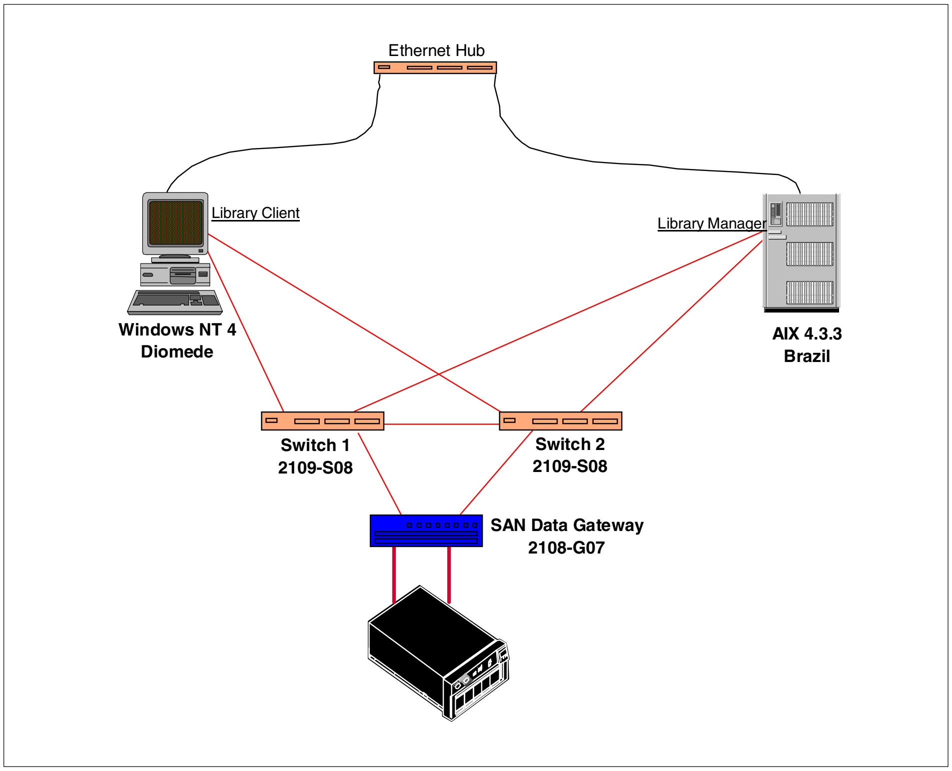

Figure 122. Multiple paths with dual HBAs

8.2.4.1 Zoning

Again, the first step is to implement zoning. In this case, we need to implement three different zoning configurations. The first one, `TSM1_Config` is used for the situation where all is normal. In that case, we connect Diomede through switch 1 to the tape library, and Brazil through switch 2. A second configuration, `TSM2_Config` is used in the case that switch 2 fails. This configuration connects Brazil using the second HBA. Finally, `TSM3_Config` is used to connect Diomede through switch 2 to the SAN data gateway. Figure 123 shows this layout.

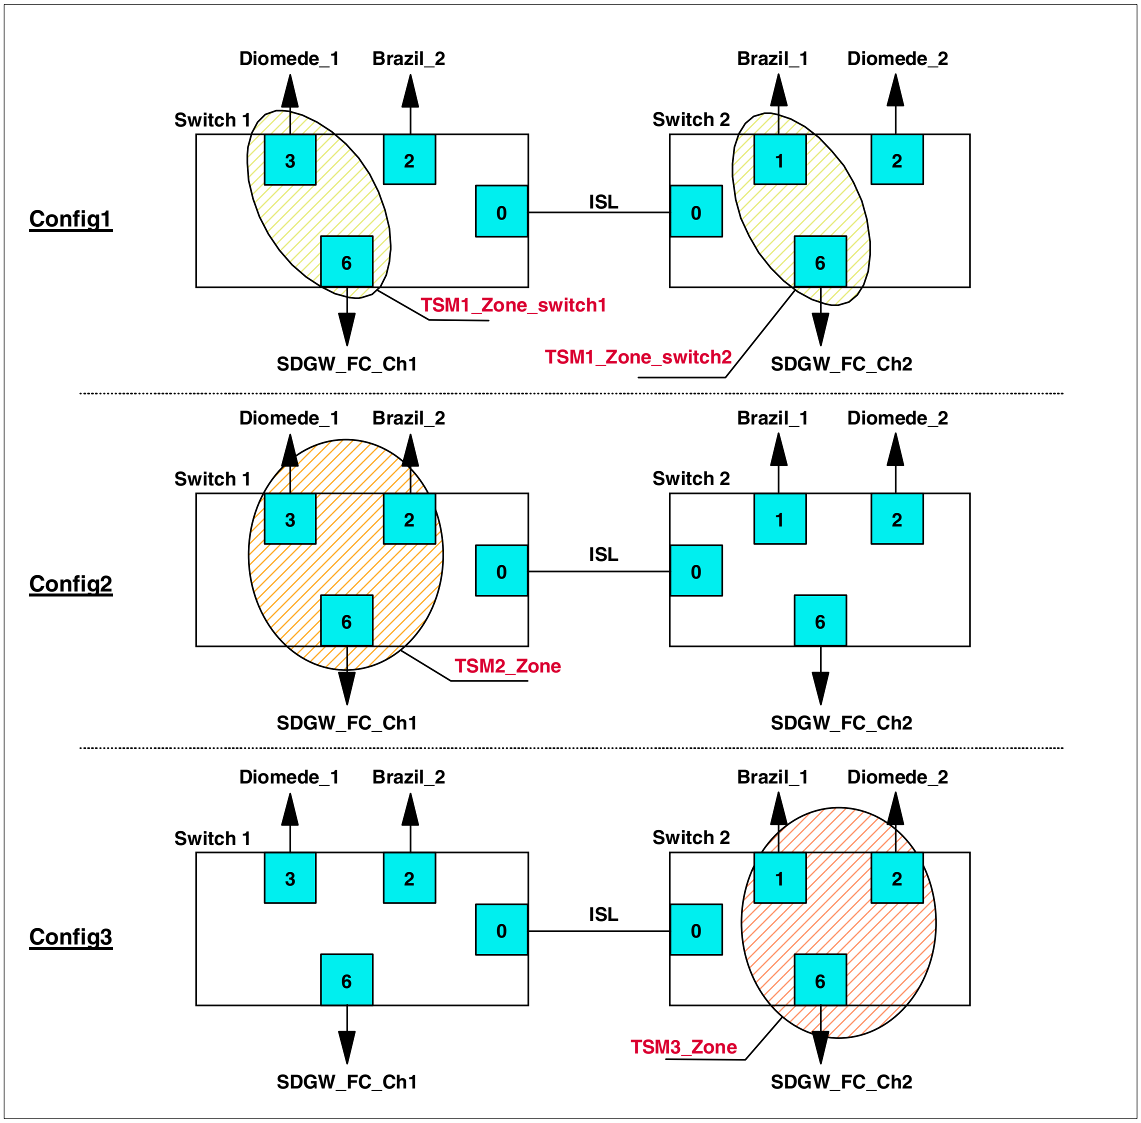

Figure 123. Zoning layout

The following shows the result of the `CfgShow` command, giving us our zoning configuration:

```
switch2:admin> cfgshow
Defined configuration:
 cfg:   Config1 TSM1_Zone_switch1; TSM1_Zone_switch2
 cfg:   Config2 TSM2_Zone
 cfg:   Config3 TSM3_Zone


zone:  TSM1_Zone_switch1
                Diomede_1; San_Data_Gateway_FC_1
zone:  TSM1_Zone_switch2
                San_Data_Gateway_FC_2; Brazil_1
zone:  TSM2_Zone
                Diomede_1; Brazil_2; San_Data_Gateway_FC_1
zone:  TSM3_Zone
                Diomede_2; San_Data_Gateway_FC_2; Brazil_1


 alias: Brazil_1                2,1
 alias: Brazil_2                1,2
 alias: Diomede_1               1,3
 alias: Diomede_2           2,2
 alias: ISL                     1,0; 2,0
 alias: Jamaica                 1,1
 alias: San_Data_Gateway_FC_1   1,6
 alias: San_Data_Gateway_FC_2   2,6

Effective configuration:
 cfg:   Config1
zone:  TSM1_Zone_switch1
                1,3
                1,6
zone:  TSM1_Zone_switch2
                2,6
                2,1
```

8.2.4.2 Recovering from a switch failure

In this case, we can look at two possible failures: one that disables switch 1, breaking the connection between the library client and the tape devices, and a second failure breaking the connection between the library server and the library device. This occurs when switch 2 fails. In this test, switch failures were simulated by disabling the switches. These tests were done during a running backup session, both on Brazil and Diomede.

Library Client connection failure

The first failure was introduced by disabling switch 1, breaking the connection between Diomede, the library client, and the tape devices.

The failure detection should be performed by the switch management system. This for the same reason as stated in the Section “Failure indications” on page 250.

The situation after the failure was as follows:

- The drive that was used by Diomede had a physical volume inserted. However, the `Query Mount` command did not show any tape mount.
- Diomede requested tape mounts for the backup session that was still running, which were satisfied by Brazil using another free device in the library. However, once mounted, Diomede was unable to detect that the volume was actually mounted in the tape device.
- The drive that was being used by Diomede when the error occurred was still physically occupied. The result of this was an I/O error on Brazil whenever it tried to use the drive. The specific message was the destination drive was full in an attempt to move a volume.

The solution resembles the one documented in the Section “Steps to recovery” on page 251. This includes the following steps:

- Disable all sessions and processes requesting tape mounts on Diomede.
- Place the failing drive offline on all Tivoli Storage Manager servers.
- Replace the switch configuration. In our example, we changed the configuration to `TSM2_Config`, enabling Diomede to connect to the library using switch 2 and its second HBA.
- Reboot all servers impacted by the new switch configuration. In this case, however, the HBA used was changed. The result was that the SCSI address changed. So, after restarting the Tivoli Storage Manager server on Diomede, we needed to update the drives, using the new SCSI address.

The previous steps led to a situation in which all Tivoli Storage Manager servers could perform tape operations, but where the library was in an inconsistent state.

To recover from that situation, halt all Tivoli Storage Manager servers, remove the tape manually from the drive, and reboot the SAN Data Gateway.

Library Manager connection failure

Finally, we tested the case in which the switch connecting the library manager, Brazil, failed. This was done when backups were running using both Brazil and Diomede as Tivoli Storage Manager servers.

The situation after the failure was as follows:

- Brazil immediately gave an I/O error on the drive.
- Diomede, however, continued with its backup process. This will continue as long as there is no need to mount or dismount a volume for Diomede.

In this situation, the impact of the failure is more severe. Since Brazil acts as our library manager, we need to get it back online, and in a consistent state. To do this, the following actions were performed:

- Begin by stopping all Tivoli Storage Manager servers.
- Remove any physical volumes left in the drive.
- Reboot the SAN Data Gateway
- Also, since the second HBA was used, the tape devices need to be re-configured. This is done using the `cgfmr` tool. Start by deleting the old devices (using the `rmdev` command), and then run `cfgmr`. This will create two new available devices.
- Restart the Tivoli Storage Manager server on Brazil and update drives and library, reflecting the new devices. When doing this, the update may not occur immediately, and the Tivoli Storage Manager server may seem to be blocked. This is due to the fact that it tries to connect to the library and the drives using the old devices. After an internal time-out, however, this situation is resolved.
- Restart the Tivoli Storage Manager server on all library clients.

8.2.5 Remarks

We would like to comment on two points that we noticed during our testing. These include the reaction of the library manager when the library client goes offline, and the problems that might occur when rejoining switches after error correction.

Library client time-out

When using a shared tape library, the library manager tests regularly if the library client is still online. The purpose of this test, that occurs every 20 minutes, is to free tape drives in case of library client failure. In our tests, however, the drive reclamation failed, due to the fact that the SCSI bus was still in use.

Below is an extract of the library manager activity log:

```
ANR9999D mmsshr.c(3071): Drive DRIVE0 in Library 3570LIB has not been confrmed for use by server
JAMAICA for over 1200 seconds - drive will be reclaimed for use by others
ANR8779E Unable to open drive /dev/rmt1, error number=16.
ANR9999D pvr.c(4443): Unable to opend device DRIVE0. Forcing dismount.
ANR8336I Verifying label of 3570 volume 085EA7 in drive DRIVE0 (/dev/rmt1).
ANR8311E An I/O error occurred while accessing drive DRIVE0 (/dev/rmt1) for REW operation,
errno = 9.
ANR8355E I/O error reading label for volume 085EA7 in drive DRIVE0 (/dev/rmt1).
ANR8311E An I/O error occurred while accessing drive DRIVE0 (/dev/rmt1) for QRYINQUIRY operation, errno = 9.
ANR8779E Unable to open drive /dev/rmt1, error number=16.
ANR8469E Dismount of 3570 volume 085EA7 from drive DRIVE0(/dev/rmt1) in library 3570LIB failed.
ANR8779E Unable to open drive /dev/rmt1, error number=16.
ANR9730I Manual intervention required for library device 'DRIVE0 (/dev/rmt1)'.
```

Another consequence of this frequent checking is that the network becomes an important factor. Since this library client confirmation check will occur using server-to-server communications, the network should be available.

Joining switches

Also worth mentioning are the problems that might occur when rejoining switches. As stated before, a switched fabric will distribute zoning configurations using the ISL. This means that in case that ISL fails, or the entire switch, the updates to the configuration are not distributed. When all connections are reestablished, however, the switches will try to consolidate their zoning configuration. This can lead to three types of error conditions:

- Zoning is enabled in both fabrics (or switches) and the zone configuration that is enabled is different. This is defined as a configuration mismatch.
- The name of a zone object in one fabric is used for a different type of type of zone object in the other fabric. This is called a type mismatch.
- The definition of a zone in one fabric is different from its definition in the other fabric, which is defined as a content mismatch.

The two errors that are most likely to occur are the configuration and the content mismatch. The first error is due to the fact that recovering after initial failure required re-configuration of the switch. The switch that was down at that moment did not receive that update, and it restarts using the old configuration. For the second error, we need to introduce the concept of the principal switch. In a fabric, there will always be one switch that acts in the role of principal switch. If the failed switch had the principal role, the fabric recovery routines detects this after failure, and names another switch as principal. When the original switch comes back online, however, it will present itself as the principal switch, thus creating a conflict.

Both errors result in a non-functional ISL. The easiest way to solve this problem is to reset the configuration of the failed switch before bringing it back online. In that case, it will receive a valid configuration from the other switches. The procedure to reset the switch is as follows:

1. Start a telnet session to the switch.
2. Run the `CfgClear` command, which clears the current switch configuration.
3. After running the clear command, it is necessary to set this cleared configuration in the switches NVRAM. This is done using the `CfgSave` command.
4. Restart the switch using the `Fastboot` command.

8.2.6 Conclusion

The examples above taught us that SAN fabric failures need manual interventions on the Tivoli Storage Manager servers, in order to recover or keep working. The following list gives some points that should be considered when using a multi-path switch fabric.

- Set up switch zoning.
- Use a SAN management tool, or at least trap the SNMP errors generated by the SAN fabric components.
- If a failure occurs, start by canceling and disabling all processes and client sessions running on the Tivoli Storage Manager servers connected through the failing component.
- In case the drive is in an open state, place it offline to prevent I/O errors.
- Reconfigure the switches, so that an alternate path to the devices is possible.
- Reboot (for Windows systems) or reconfigure (for UNIX systems), all of the systems that are affected by the new switch configuration.
- Plan a halt of all Tivoli Storage Manager servers, after which the failing drives can be manually unmounted, and the SAN Data Gateway rebooted to reset the SCSI bus.

Appendix A. Configuring StorWatch Specialist tools

IBM's StorWatch Specialist tools provide an easy, graphical way to manage your IBM SAN products.

This appendix covers configuration and use of these tools for the SAN components used in this publication.

A.1 StorWatch SAN Data Gateway Specialist

This section covers the StorWatch SAN Data Gateway Specialist product for Windows. Versions of this product are also available for AIX and Sun Solaris.

The IBM StorWatch SAN Data Gateway Specialist supports the following products:

- IBM SAN Data Gateway, model 2108-G07
- IBM SAN Data Gateway Router, model 2108-R03

In this section, when we refer to “gateway”, you can assume that we refer to either device.

A.1.1 Preparation

The StorWatch SAN Data Gateway Specialist communicates with SAN devices primarily by using TCP/IP. Accordingly, its is necessary to do some minimal configuration of the gateway device before it can be managed by the StorWatch SAN Data Gateway Specialist software.

1. Begin by readying either an ASCII terminal, or a terminal emulation program running on a PC to talk to the service port (serial port) on the gateway. Set your communications for 19200 baud, 8 data bits, no parity, and xon/xoff flow control. A standard null modem DB-9 (female) to DB-9 (female) serial cable is necessary to communicate through the service port. Further details regarding service port communication cabling and settings are contained in the *Installation and User's Guide 2108 Model G07*, SC26-7304.

   When ready, apply power and turn on the gateway. The gateway should begin to display boot messages on the screen. The gateway uses a VxWorks real-time operating system, and provides a command line interface, as well as “out-band” network connectivity through the integrated ethernet port.

When bootup is complete, you should see a display similar to the following:

```
SAN Gateway Version 0340.16 Built May 19 2000, 08:27:01
Gateway
0xc1f9e2b0 (tShell): VPS Feature NOT Enabled
Data Mover is Disabled
SCSI 3 - DE - Terminated
SCSI 1 - DE - Terminated
SCSI 4 - DE - Terminated
SCSI 2 - DE - Terminated
sdgw1
IP Address set to 193.1.1.25:ffffff00
value = 0 = 0x0


Done executing startup script ffs0:sna.rc
- Service Port Reflected Signal Test
Starting shell
Gateway >
```

2. The name of the gateway cannot be set from the StorWatch GUI. Set the name of the gateway now with the `hostNameSet` command:

```
Gateway > hostNameSet "Lab_SDGW1"
Write complete
Target hostname set to Lab_SDGW1
value = 0 = 0x0
Lab_SDGW1 >
```

3. Set the IP address and network mask for the ethernet port using the `ethAddrSet` command:

```
Lab_SDGW1 > ethAddrSet "193.1.1.25","255.255.255.0"
Inet Mask set to ffffff00 for Ethernet interface
Write complete
Host Address set to 193.1.1.25 for Ethernet interface
value = 0 = 0x0
```

4. Verify the address using the `ifShow` command:

```
Lab_SDGW1 > ifShow
lnPci (unit number 0):
     Flags: (0x63) UP BROADCAST ARP RUNNING
     Internet address: 193.1.1.25
     Broadcast address: 193.1.1.255
     Netmask 0xffffff00 Subnetmask 0xffffff00
     Ethernet address is 00:60:45:16:0b:a3
     Metric is 0
     Maximum Transfer Unit size is 1500
     24 packets received; 2 packets sent
     0 input errors; 0 output errors
     0 collisions
lo (unit number 0):
     Flags: (0x69) UP LOOPBACK ARP RUNNING
     Internet address: 127.0.0.1
     Netmask 0xff000000 Subnetmask 0xff000000
     Metric is 0
     Maximum Transfer Unit size is 4096
     6 packets received; 6 packets sent
     0 input errors; 0 output errors
     0 collisions
value = 18 = 0x12
```

5. After setting the parameters listed above, reboot the gateway using the `reboot` command:

```
Lab_SDGW1 > reboot
==  reboot
c00dc048 - _reboot
c011f41c - _yystart
c0120f54 - _yyparse
c00dad0c - _execute
c00daba8 - _shell
c00da9a8 - _shell
c011876c - _vxTaskEntry
VCM: 20 tasks started
MGT List Count 0
CMD List Count 0
LBuffer Pools:
 2097152 bytes: Free   8 of   8 - Wait   0,  0, Sig   0 N   0 L   0 list 0xc01e
  262144 bytes: Free 128 of 128 - Wait   0,  0, Sig   0 N   0 L   0 list 0xc01e
   32768 bytes: Free 256 of 256 - Wait   0,  0, Sig   0 N 227 L 227 list 0xc01e
    4096 bytes: Free 648 of 1664 - Wait   0,  0, Sig   0 N 1038 L  22 list 0xc0

Syncing volume ffs0. Volume now unmounted.
Flash File Systems now synced.
```

At this point, the gateway is ready to be attached to the ethernet network. Attach the cable to the network, and verify that the server which will be running the StorWatch software can “see” the gateway using `ping` and `telnet` commands.

A.1.2 Installation

StorWatch SAN Data Gateway Specialist is available for free download from:

`http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/NT/sdgnt.htm`

StorWatch SAN Data Gateway Specialist for Windows downloads as an executable file. Download the product to an empty directory, and start the installer by double-clicking on the downloaded file.

Installation uses the standard Windows installer. Normally, the defaults can be used, with the exception that if you do not have a StorWatch server located somewhere on your network, you will need to install that as well. The default is to install only the client component.

If you do **not** have a StorWatch server installed elsewhere on your network, select the check box to install the StorWatch server when the dialog box shown in Figure 124 appears.

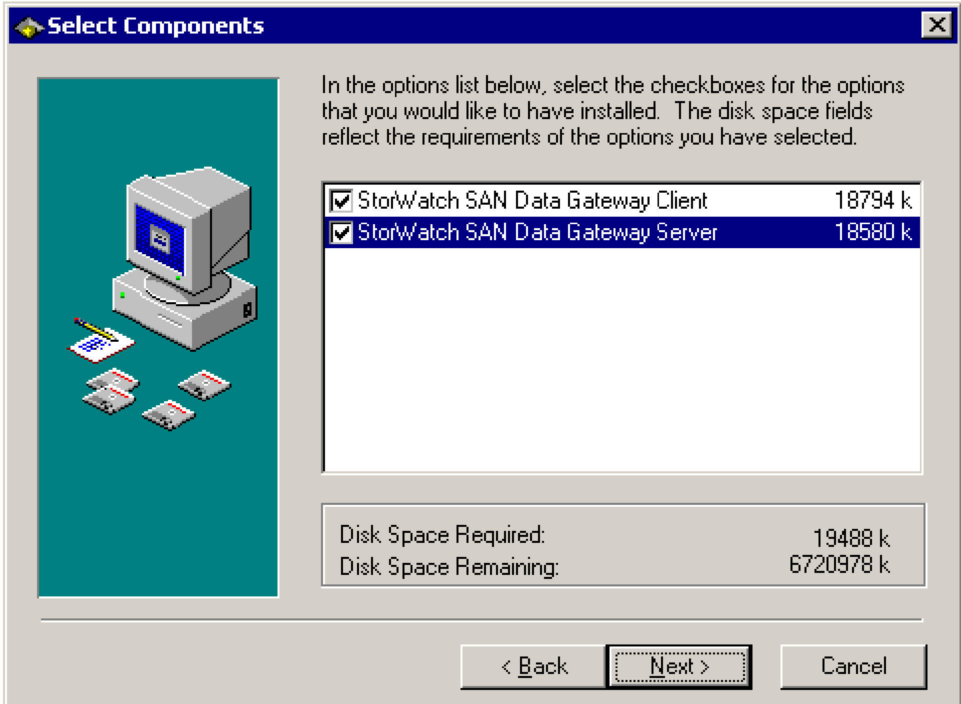

Figure 124. Install client and server

After the installation routine completes, you *must* reboot in order to properly initialize the StorWatch server component.

A.1.3 Operation of the SAN Data Gateway Specialist

After installing the StorWatch SAN Data Gateway Specialist and rebooting, you start the StorWatch server. Click on **Start -> Programs -> IBM StorWatch SAN Data Gateway Specialist -> Server**. This will open an `iconified` command window, which will reside in the Windows task bar; it will appear as shown in Figure 125.

Figure 125. StorWatch server iconified

Interaction with the StorWatch server is only necessary when shutting down the server; we found that opening the server window and typing “exit” or “quit” brings the StorWatch server down more cleanly.

Once the server is running, start the StorWatch client by clicking on **Start -> Programs -> IBM StorWatch SAN Data Gateway Specialist -> Client**. You will be presented with the logon screen shown in Figure 126.

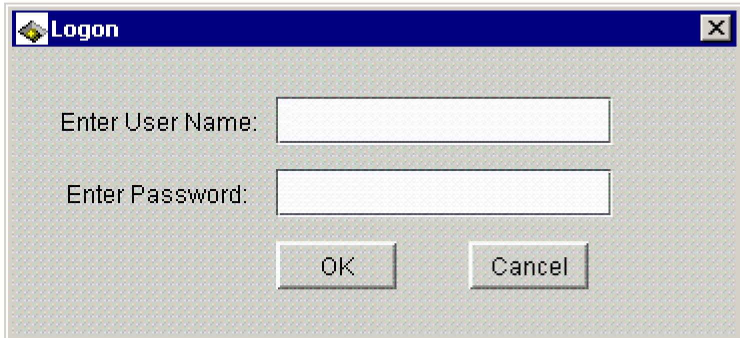

Figure 126. Client logon

Initially, the user name will be “StorWatch”, with a password of “StorWatch” (case sensitive). After a successful logon, you may see a window asking you which device to attach to. When the StorWatch server starts, it scans the SAN looking for devices that it can manage. The list you see in this window is the list of devices that the server discovered. Select the desired device from the pull-down menu, as shown in Figure 127.

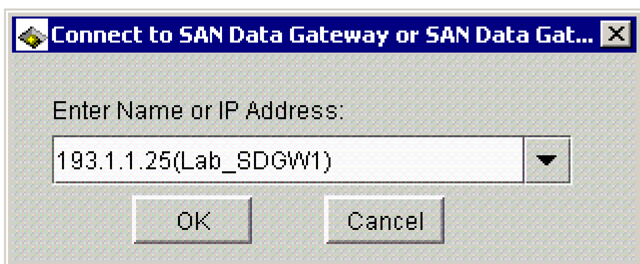

Figure 127. Device connection dialog

After selecting the device, you will see the main window, shown in Figure 128.

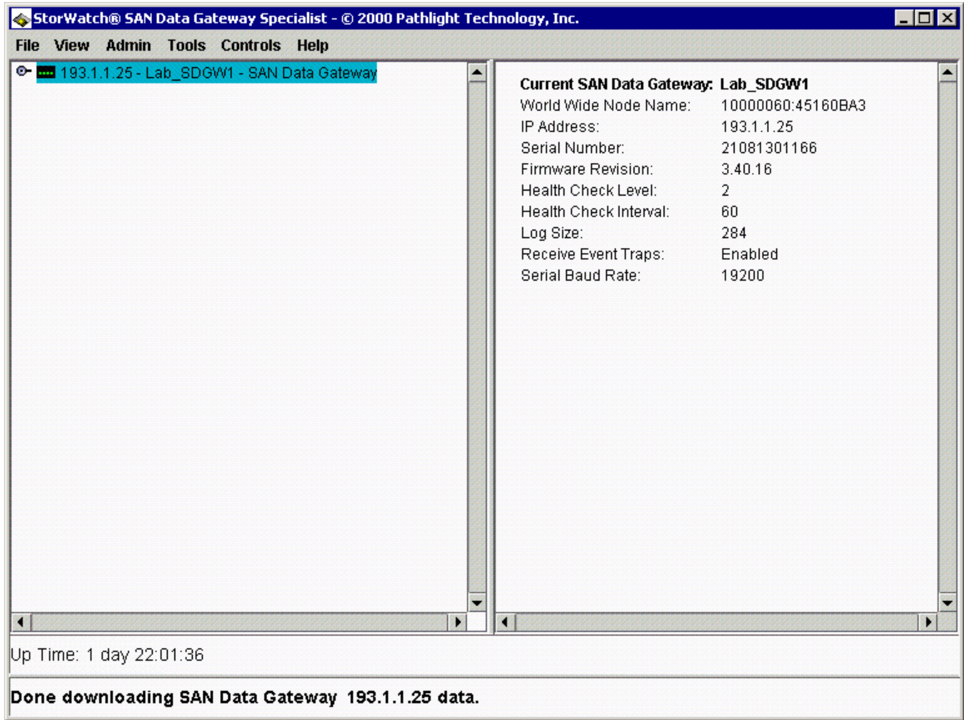

Figure 128. Client main window

By clicking on the small bullet in the left pane next to the gateway device, you will get a representation of the devices attached to the gateway, as shown in Figure 129.

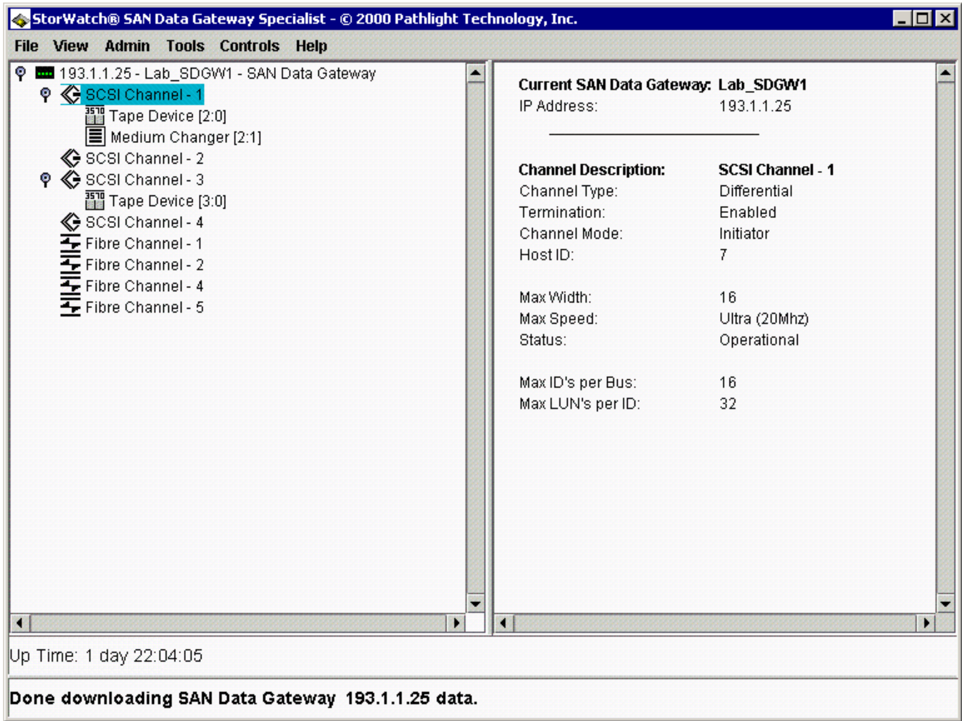

Figure 129. Devices attached to the gateway

In addition to the main window, you will also get an Event Traps window, as shown in Figure 130.

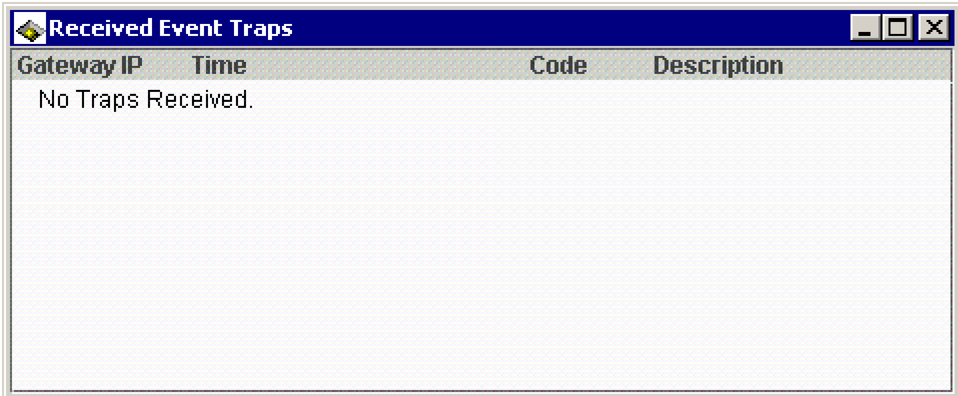

Figure 130. Event Traps window

At this point, you have completed the installation of the SAN Data Gateway Specialist. Refer to the publication *Installation and User's Guide 2108 Model G07*, SC26-7304 for further information on the use of this tool, including information on channel zoning and defining virtual private SANs.

A.2 StorWatch Fibre Channel Switch Specialist

The IBM StorWatch Fibre Channel Switch Specialist provides an easy-to-use graphical interface by which you can manage IBM Fibre Channel Switches and Managed Hubs. IBM products which can be managed by this Specialist are:

- IBM 3534 SAN Fibre Channel Managed Hub
- IBM 2109-S08 SAN Fibre Channel Switch
- IBM 2109-S016 SAN Fibre Channel Switch

In this section, when we say "switch", you can assume that we refer to any of the above listed products.

Unlike the StorWatch Specialist for the SAN Data Gateway and Router product, it does not require you to download and install a regular Windows-based program; instead, this Specialist runs as a Java applet in your Web browser.

This Java based utility requires the Java Runtime Environment version 1.2.2. The JRE can be downloaded free of charge from:

`http://www.javasoft.com/products/`

The download page for the Fibre Channel Switch is located at:

`http://www.storage.ibm.com/hardsoft/products/fcswitch/download.htm`

As of this writing, the supported platforms for the StorWatch Fibre Channel Switch Specialist are Windows and Solaris.

After downloading the Java Runtime Environment, begin the installation by double-clicking on the installation file. The installation routine utilizes the standard Windows installer. We suggest that you reboot after installation, even though the installation routine does not require that you do so.

Once you have installed the JRE and rebooted, launch your browser, and direct it at the IP address of the switch that you wish to manage. After the applet loads, you should see a display similar to Figure 131.

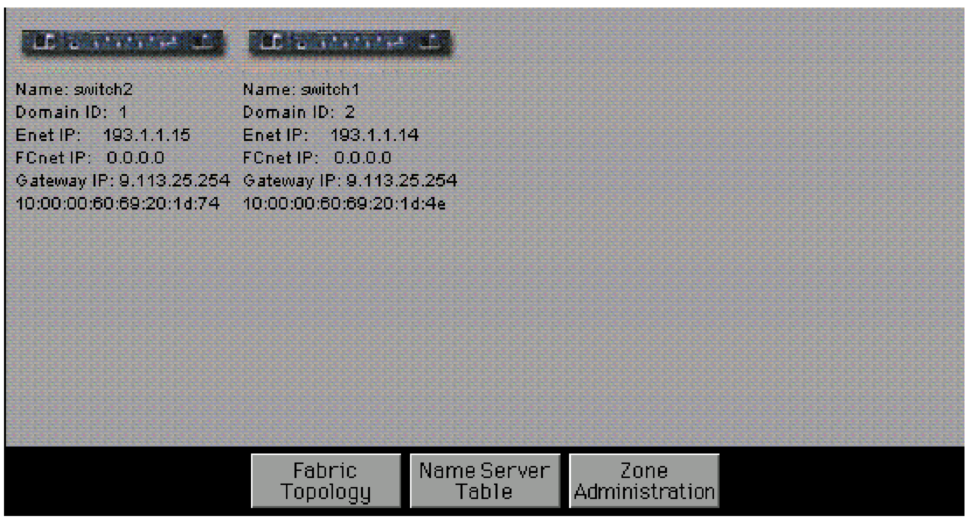

Figure 131. Switch fabric display

In Figure 131, selecting the **Zone Administration** icon will bring up a window that allows you to define your fabric configuration, including aliases and zoning.

Clicking on one of the switches in the display in Figure 131 will bring up an animated graphic of the switch front panel, as shown in Figure 132:

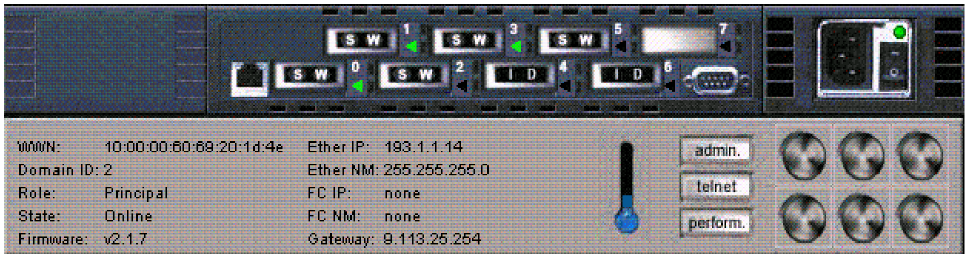

Figure 132. Main switch display

This view provides access to several items useful for managing an individual switch, including a telnet window and reporting statistics for individual ports.

Design of a SAN fabric and detailed information regarding configuration of Fibre Channel switches is beyond the scope of this publication; refer to the installation and user's guides for the particular model of switch you are using.

Other sources of information include:

- *Introduction to Storage Area Network*, SG14-5470
- *Tape Future in Fabrics*, SG24-5474

Appendix B. Configuring SLIC Manager on Windows

Basic configuration of the Vicom SLIC Router hardware is covered in Section 3.2.5, “Vicom SLIC Router FC-SL basic configuration” on page 72. The configuration information provided in this section details the steps necessary to initialize the router, and to map the physical drives to the router such that host systems can access them. This appendix has been included as an aid to installation and configuration of Vicom System’s SLIC Manager utility for their SLIC Router products.

The SLIC Manager software is provided on a CD that ships with the SLIC Router, and is not, at this time, available for download from the Vicom Web site.

We advise you to refer to the installation and user documentation, which is provided with the Vicom SLIC Router.

B.1 Loading the software

Note:

The steps in this section make the assumption that you have read and completed the Vicom SLIC Router setup instructions provided in Section 3.2.5, “Vicom SLIC Router FC-SL basic configuration” on page 72.

Insert the CD, click on **Start**, then **Run**. Select the file shown in Figure 133.

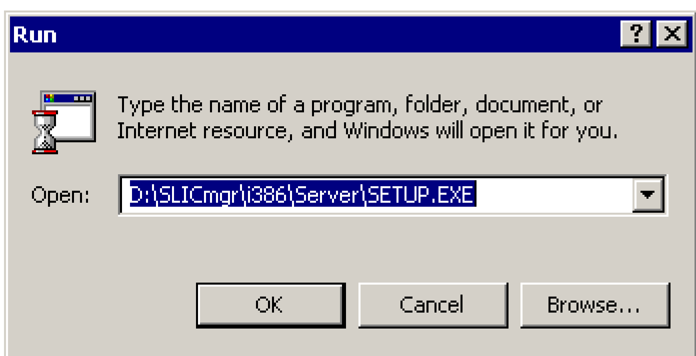

Figure 133. SLIC Manager install path

The installation process uses the standard Windows installer. During the installation routine, you will be notified that both the server and client components will be installed, as shown in Figure 134.

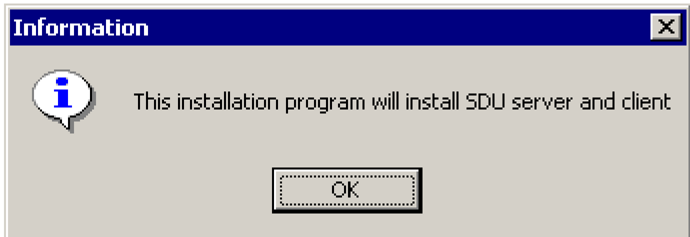

Figure 134. Installation of server and client

We did not find it necessary to reboot after completing the installation.

B.2 Setting up the slicd daemon configuration file

After installation, it is necessary to set up the `slicd` daemon configuration file, which is located in the SLIC Manager server directory (the default is `c:\ibm7190\sdus`).

With your favorite text editor, open the file `7190.CFG`.

The default file has all lines commented out; you are interested in the following stanza:

```
# Example:
# c0 = {
#       path = J:\IBM7190.SFA, K:\IBM7190.SFA;
# # QueryChangeInterval in seconds
#       QueryChangeInterval=10;
# };
```

In the example text, `c0` is the name by which the SLIC Router will be known to the system. You may change to any name you choose, but the default will work as-is. The “`path=`” statement must be changed to point to a drive letter that is on one of the SSA disks attached to the router, and which can be accessed by this system. At least one file must be included in the path statement, although more (on different disks) are recommended.

The completed statement for our lab system follows:

```
# Example:
 c0 = {
       path = e:\IBM7190.SFA;
# QueryChangeInterval in seconds
       QueryChangeInterval=10;
 };
```

After these modifications, save the file, making sure that it is saved as a `.cfg` file, and not as a `.txt` file.

B.3 Test the slicd daemon

Open a DOS window, and change to the directory where the `slicd` daemon resides; this will typically be `c:\ibm7190\sdus`. Start the daemon in debug mode to test your configuration:

```
C:\ibm7190\sdus>slicd -debug
Debugging Vicom SLIC Daemon service.

Vicom SLIC Manager -- SLIC Daemon Version 1.0.7
Copyright (C) 1999, VICOM SYSTEMS Inc. All rights reserved.
Waiting for Network Resources.

Query Change Application Started
Health Checking...
Completed First Health Check.
```

If your configuration is valid, your output should appear as shown. Exit the debug mode by pressing “`control-c`”

Install the `slicd` daemon using this command:

```
C:\ibm7190\sdus>slicd -install
Vicom SLIC Daemon service installed.
```

You should now be able to start the SLIC daemon as a service. Figure 135 shows a screen shot from a Windows 2000 server:

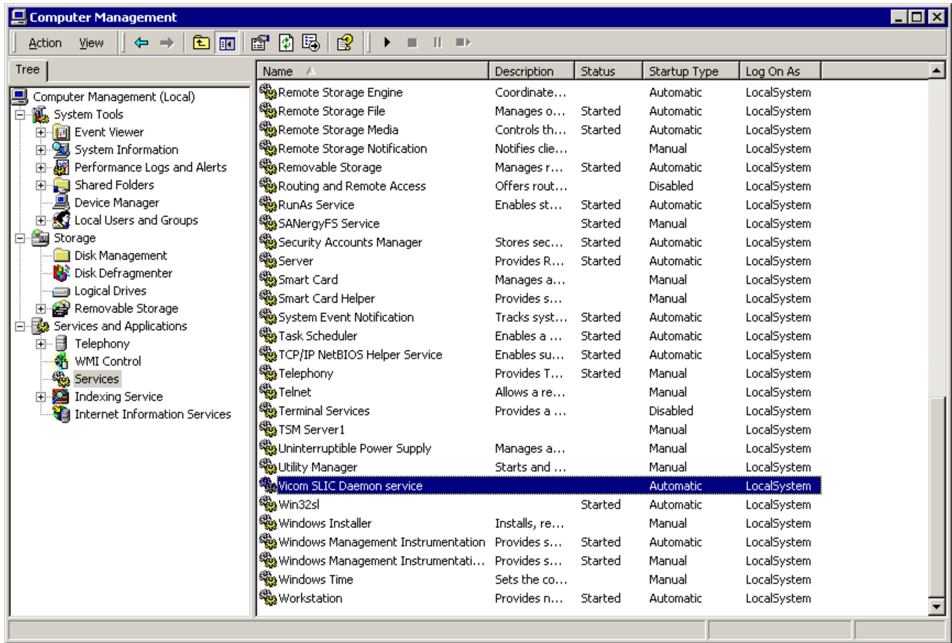

Figure 135. Services window showing SLIC daemon

Once the SLIC daemon is running, you can communicate with the SLIC Router using either the SLIC Manager GUI or command line.

B.4 SLIC Manager GUI introduction

Launch the GUI by clicking **Start** - **Programs** - **Vicom** - **SLIC Manager**.

You will be prompted to provide the server that has the SLIC daemon running on it, and to specify the router you wish to connect to, as in Figure 136:

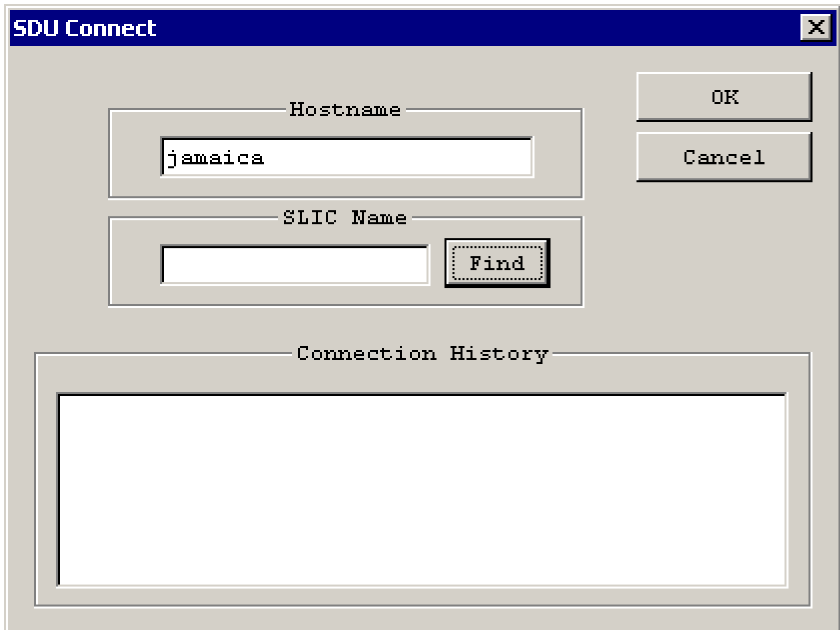

Figure 136. SDU Connect window

The host name is the server that has the SLIC daemon running on it; in our example, it is “Jamaica”, the same server that the client GUI is located on.

The SLIC name is the name that you define for the router in the SLIC daemon configuration file; we used the default of “c0”. Clicking the **Find** button will present a list of routers present in the configuration file, as in Figure 137.

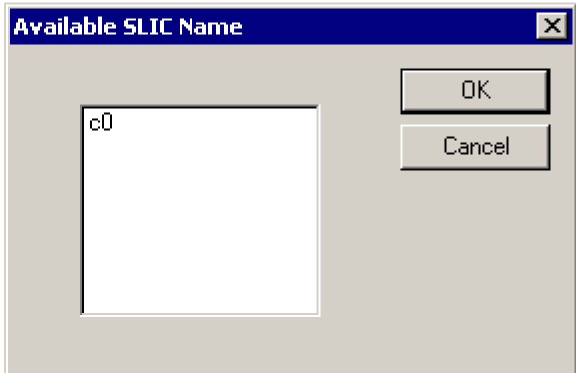

Figure 137. SLIC names available

After providing your selections, you should see the (rather sparse) main screen. All of the management functions are accessed via the drop-down menus.

Refer to the SLIC Router product documentation for further information on using the GUI.

B.5 SLIC Manager command line introduction

The SLIC Manager product provides a number of command line utilities, all of which are located in the client directory, typically `c:\ibm7190\sduc`. Documentation of the use of these utilities is covered in the SLIC Manager documentation.

A couple of these utilities we found to be of particular use:

ssa_serv
The ssa_serv utility can be started with the command:

c:\ibm7190\sduc\ssa_serv -dc0

Where the trailing “c0” is the name defined in the `slicd` configuration file.

You should then be greeted with the following menu:

```
C:\ibm7190\sduc>ssa_serv -dc0

Vicom SLIC Manager -- SSA_SERV Version 1.0.7
Copyright (C) 1999, VICOM SYSTEMS Inc. All rights reserved.

You may not modify, reverse engineer, decompile, or otherwise attempt to derive
the source code for this Software. For more information regarding Your rights
to use and copy this Software, consult Your user's manual.


1. Download Microcode
2. Display VPD
3. Identify a Disk Drive Module
4. Set/Reset Service mode
5. Disk Drive Diagnostics
6. SLIC Diagnostics
7. Error Log Analysis
8. Clear Check Mode
9. Activity Monitor
a. Enclosure Service Aid
b. Clean Up Configuration Table
c. Upload/Download Configuration
d. Show Fibre Channel Map
m. Force SSA Master
p. Print Fibre Channel Map and SSA Topology
t. Show SSA Topology
q. Quit
?. Menu
>>
```

Most of the options displayed in ssa_serv are self explanatory; refer to the product documentation if further information is required.

showmap

Showmap provides a quick snapshot of the SLIC Router's configuration, which we found quite useful. The output is as follows:

```
C:\ibm7190\sduc>showmap -dc0

*** SSA TOPOLOGY ***

I00000.c0.00000060-22003875.N.OK.FCSL-H***

SSA.PORT.A1=
=T00007.00000020-35E684AE.N.OK.SSA-DRIVE=
=T00006.00000020-35E68462.N.OK.SSA-DRIVE=
=T00005.00000020-35E6840E.N.OK.SSA-DRIVE=
=T00004.00000020-35E22D6C.N.OK.SSA-DRIVE=
=T00003.00000020-35E6848D.N.OK.SSA-DRIVE=
=T00002.00000020-35E69318.N.OK.SSA-DRIVE=
=T00001.00000020-35E22D0B.N.OK.SSA-DRIVE=
=T00000.00000020-35E68229.N.OK.SSA-DRIVE=
=SSA.PORT.A2

OFFLINE:

Storage Capacity for Each Physical Drive:
Target Number          Device Name     Capacity (MegaByte)

T00000                 DRVC0016           8689 MB
T00001                 DRVC0016           8689 MB
T00002                 DRVC0016           8689 MB
T00003                 DRVC0016           8689 MB
T00004                 DRVC0016           8689 MB
T00005                 DRVC0016           8689 MB
T00006                 DRVC0016           8689 MB
T00007                 DRVC0016           8689 MB

Properties of Logical Drives:
Target Number          Device Name     Capacity
T33025                      lad          17378 MB

Fibre Channel LUN Map:
FC      FC      Mirror  InstCpy Composite       Spare   Target
ID      LUN     Drive   Drive   Drive                   Number
005     000     Null    Null    Null            Null    T00000
005     001     Null    Null    Null            Null    T00001
005     002     Null    Null    Null            Null    T00002
005     003     Null    Null    Null            Null    T00003
005     004     Null    Null    Null            Null    T00004
005     005     Null    Null    Null            Null    T00005
005     009     Null    Null    T33025          Null    T00006
005     009     Null    Null    T33025          Null    T00007

Private Drives and Their Owner:
Volume ID       Owner
```

Appendix C. Special notices

This publication is intended to help IBM, Tivoli, customer, vendor, and consulting personnel who wish either to gain an understanding of how SAN technology applies to a Tivoli Storage Manager environment, or to obtain assistance in implementing Tivoli Storage Manager utilizing SAN-attached storage. The information in this publication is not intended as the specification of any programming interfaces that are provided by Tivoli Storage Manager and Tivoli SANergy File Sharing. See the PUBLICATIONS section of the IBM Programming Announcement for Tivoli Storage Manager and Tivoli SANergy File Sharing for more information about what publications are considered to be product documentation.

References in this publication to IBM products, programs or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program or service.

Information in this book was developed in conjunction with use of the equipment specified, and is limited in application to those specific hardware and software products and levels.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785.

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, Dept. 600A, Mail Drop 1329, Somers, NY 10589 USA.

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

The information contained in this document has not been submitted to any formal IBM test and is distributed AS IS. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

Any pointers in this publication to external Web sites are provided for convenience only and do not in any manner serve as an endorsement of these Web sites.

The following terms are trademarks of the International Business Machines Corporation in the United States and/or other countries:

IBM ®
S/390
SP
StorWatch
System/390
VM/ESA
Wave
Wizard
Redbooks
Redbooks Logo
XT
400
Lotus
Lotus Notes
Domino
Notes

The following terms are trademarks of other companies:

Tivoli, Manage. Anything. Anywhere.,The Power To Manage., Anything. Anywhere.,TME, NetView, Cross-Site, Tivoli Ready, Tivoli Certified, Planet Tivoli, and Tivoli Enterprise are trademarks or registered trademarks of Tivoli Systems Inc., an IBM company, in the United States, other countries, or both. In Denmark, Tivoli is a trademark licensed from Kjøbenhavns Sommer - Tivoli A/S.

C-bus is a trademark of Corollary, Inc. in the United States and/or other countries.

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

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States and/or other countries.

PC Direct is a trademark of Ziff Communications Company in the United States and/or other countries and is used by IBM Corporation under license.

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States and/or other countries.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Other company, product, and service names may be trademarks or service marks of others.

Appendix D. Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

D.1 IBM Redbooks

For information on ordering these publications see "How to get IBM Redbooks" on page 289.

- *Introduction to Storage Area Network, SAN*, SG24-5470
- *Storage Area Networks: Tape Future in Fabrics*, SG24-5474
- *IBM Magstar Tape Products Family: A Practical Guide*, SG24-4632
- *Guide to Sharing and Partitioning IBM Tape Library Dataservers*, SG24-4409
- *Using ADSM to Back Up and Recover Microsoft Exchange Server*, SG24-5266
- *R/3 Data Management Techniques Using Tivoli Storage Manager,* SG24-5743
- *Tivoli Storage Manager 3.7.3 & 4.1 Technical Guide,* SG24-6110

D.2 IBM Redbooks collections

Redbooks are also available on the following CD-ROMs. Click the CD-ROMs button at ibm.com/redbooks for information about all the CD-ROMs offered, updates and formats.

CD-ROM Title	Collection Kit Number
IBM System/390 Redbooks Collection	SK2T-2177
IBM Networking Redbooks Collection	SK2T-6022
IBM Transaction Processing and Data Management Redbooks Collection	SK2T-8038
IBM Lotus Redbooks Collection	SK2T-8039
Tivoli Redbooks Collection	SK2T-8044
IBM AS/400 Redbooks Collection	SK2T-2849
IBM Netfinity Hardware and Software Redbooks Collection	SK2T-8046
IBM RS/6000 Redbooks Collection	SK2T-8043
IBM Application Development Redbooks Collection	SK2T-8037
IBM Enterprise Storage and Systems Management Solutions	SK3T-3694

D.3 Other resources

These publications are also relevant as further information sources:

- *Fibre Channel: Connection to the Future,* ISBN 1-878701-45-0
- *Installation and User's Guide 2108 Model R03*, SC26-7355
- *Installation and User's Guide 2108 Model G07*, SC26-7304
- *IBM SAN Fibre Channel Switch 2109 Model S08 Users Guide,* SC26-7349
- *Magstar 3494 Tape Library Introduction and Planning Guide*, GA32-0279
- *Magstar 3494 Tape Library Operator Guide*, GA32-0280
- *Tivoli Storage Manager for Windows Administrator's Guid*e, GC35-0410
- *Tivoli Storage Manager for Windows Administrator's Reference*, GC35-0411
- *Tivoli Storage Manager for AIX Administrator's Guide*, GC35-0403
- *Tivoli Storage Manager for AIX Administrator's Reference*, GC35-0404
- *ADSM V3R1 Using the Application Program Interface*, SH26-4081

D.4 Referenced Web sites

These Web sites are also relevant as further information sources:

- `http://www.storage.ibm.com/ibmsan/products/sanfabric.htm`
  IBM's Enterprise Storage Area Network main site
- `http://www.storage.ibm.com/hardsoft/products/fcswitch/supserver.htm`
  Supported server information for IBM SAN Fibre Channel Switch 2109-S08
- `http://www.storage.ibm.com/hardsoft/products/fcswitch/download.htm`
  Software and firmware download site for IBM SAN Fibre Channel Switch 2109-S08
- `http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/AIX/sdgaix.htm`
  Software and firmware download site for IBM SAN Data Gateway 2108-G07 for AIX
- `http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/NT/sdgnt.htm`
  Software and firmware download site for IBM SAN Data Gateway 2108-G07 for Windows NT and Windows 2000

- http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/Solaris/sdgsun.htm
  Software and firmware download site for IBM SAN Data Gateway 2108-G07 for Sun Solaris
- http://www.storage.ibm.com/hardsoft/products/sangateway/supserver.htm
  Supported server information for 2108-G07 SAN Data Gateway
- http://www.storage.ibm.com/hardsoft/products/sangateway/support/cdr/sdgcdr.htm
  Software and firmware download site for IBM SAN Data Gateway Router 2108-R03
- http://www.storage.ibm.com/hardsoft/products/fchub/fcshub.htm
  Product information for 2107-H07 Fibre Channel hub
- http://www.storage.ibm.com/hardsoft/products/tape/ro3superserver.htm
  Supported server information for 2108-R03 SAN Data Gateway Router
- http://www.storage.ibm.com/hardsoft/products/tape/tapesupport.htm
  SAN Data Gateway 2108-G07 Support Matrix for Tape
- http://www.vicom.com/product_fcsl.html
  Supported Server information for Vicom Fibre Channel SLIC Router Model FC-SL
- http://www.tivoli.com/support/storage_mgr/requirements.html
  Tivoli Storage Manager product requirements and supported devices
- http://www.tivoli.com/support/storage_mgr/san/libsharing.html
  Tivoli Storage Manager Library Sharing support information
- http://www.rs6000.ibm.com/support/micro/download.html
  RS/6000 Microcode download site
- ftp://index.storsys.ibm.com/devdrvr/
  IBM SSD Open Systems Device Driver FTP site
- ftp://index.storsys.ibm.com/tivoli-storage-management/maintenance/server/v4r1/
  Tivoli Storage Manager Release 4.1 maintenance update FTP site
- http://www.tivoli.com/products/solutions/storage/
  Tivoli Storage Management home page
- http://www.tivoli.com/products/index/storage_mgr/
  Tivoli Storage Manager home page
- http://www.qlogic.com/bbs-html/csg_web/adapter_pages/fc_adapters.html
  QLogic Fibre Channel adapters — supported operating systems
- http://www.javasoft.com/products
  JAVA technology, products, and APIs

How to get IBM Redbooks

This section explains how both customers and IBM employees can find out about IBM Redbooks, redpieces, and CD-ROMs. A form for ordering books and CD-ROMs by fax or e-mail is also provided.

- **Redbooks Web Site** **ibm.com**/redbooks

  Search for, view, download, or order hardcopy/CD-ROM Redbooks from the Redbooks Web site. Also read redpieces and download additional materials (code samples or diskette/CD-ROM images) from this Redbooks site.

  Redpieces are Redbooks in progress; not all Redbooks become redpieces and sometimes just a few chapters will be published this way. The intent is to get the information out much quicker than the formal publishing process allows.

- **E-mail Orders**

  Send orders by e-mail including information from the IBM Redbooks fax order form to:

	e-mail address
In United States or Canada	pubscan@us.ibm.com
Outside North America	Contact information is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Telephone Orders**

United States (toll free)	1-800-879-2755
Canada (toll free)	1-800-IBM-4YOU
Outside North America	Country coordinator phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

- **Fax Orders**

United States (toll free)	1-800-445-9269
Canada	1-403-267-4455
Outside North America	Fax phone number is in the “How to Order” section at this site: http://www.elink.ibmlink.ibm.com/pbl/pbl

This information was current at the time of publication, but is continually subject to change. The latest information may be found at the Redbooks Web site.

IBM Intranet for Employees

IBM employees may register for information on workshops, residencies, and Redbooks by accessing the IBM Intranet Web site at http://w3.itso.ibm.com/ and clicking the ITSO Mailing List button. Look in the Materials repository for workshops, presentations, papers, and Web pages developed and written by the ITSO technical professionals; click the Additional Materials button. Employees may access MyNews at http://w3.ibm.com/ for redbook, residency, and workshop announcements.

IBM Redbooks fax order form

Please send me the following:

Title	**Order Number**	**Quantity**

First name Last name

Company

Address

City Postal code Country

Telephone number Telefax number VAT number

☐ Invoice to customer number

☐ Credit card number

Credit card expiration date Card issued to Signature

We accept American Express, Diners, Eurocard, Master Card, and Visa. Payment by credit card not available in all countries. Signature mandatory for credit card payment.

Index

Numerics

A

B

C

D

T

U

V

W

Z

IBM Redbooks review

Your feedback is valued by the Redbook authors. In particular we are interested in situations where a Redbook "made the difference" in a task or problem you encountered. Using one of the following methods, **please review the Redbook, addressing value, subject matter, structure, depth and quality as appropriate.**

- Use the online **Contact us** review redbook form found at **ibm.com**/redbooks
- Fax this form to: USA International Access Code + 1 845 432 8264
- Send your comments in an Internet note to redbook@us.ibm.com

Document Number **Redbook Title**	SG24-6132-00 Using Tivoli Storage Manager in a SAN Environment
Review	
What other subjects would you like to see IBM Redbooks address?	
Please rate your overall satisfaction:	O Very Good O Good O Average O Poor
Please identify yourself as belonging to one of the following groups:	O Customer O Business Partner O Solution Developer O IBM, Lotus or Tivoli Employee O None of the above
Your email address: The data you provide here may be used to provide you with information from IBM or our business partners about our products, services or activities.	 O Please do not use the information collected here for future marketing or promotional contacts or other communications beyond the scope of this transaction.
Questions about IBM's privacy policy?	The following link explains how we protect your personal information. **ibm.com**/privacy/yourprivacy/

IBM

Redbooks

Using Tivoli Storage Manager in a SAN Environment

(0.5” spine)
0.475”<->0.875”
250 <-> 459 pages

Using Tivoli Storage Manager in a SAN Environment

Tape library sharing among multiple Tivoli Storage Manager servers

LAN-free backup of application data using the SAN

Configuring a SAN for backup purposes

Storage Area Networks (SANs) are changing the way data is moved, stored and managed, and stand to transform the computer landscape in much the same way that LANs did in the 1980s. Tivoli Storage Manager has been enhanced in the latest releases to provide new functionality that utilizes SAN technology.

This IBM Redbook will help the reader to gain a broad understanding of SANs. It provides practical implementation details for using Tivoli Storage Manager products to tap the new speed, flexibility, and manageability afforded by a SAN.

This book is written for IBM, Tivoli, customer, vendor, and consulting personnel who wish to gain an understanding of how SAN technology applies to a Tivoli Storage Manager environment, or to obtain assistance in implementing Tivoli Storage Manager utilizing SAN-attached storage.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks

SG24-6132-00 ISBN 0738421197